

Are causes ever too strong?

Downward monotonicity in the causal domain

Dean McHugh*

Institute for Logic, Language and Computation
University of Amsterdam

Abstract. Is the truth of a causal claim always preserved by strengthening the cause? For instance, does “Alice flicking the switch caused the light to turn on” entail “Alice flicking the switch and it raining in New Zealand caused the light to turn on”? We argue for this entailment, proposing that causal claims are downward monotone in their cause: if C^+ entails C then $(C \text{ caused } E)$ entails $(C^+ \text{ caused } E)$. In other words, causes are never too strong. We argue for this by presenting examples of causal claims that are assertable even though the cause is stronger than required for the claim to be true (§2). These data challenge accounts (the most prominent of which is Halpern, *Actual Causality* 2016) that predict such sentences to be false. Instead, we trace differences in their acceptability to their scalar implicatures (§3). Finally, we show that Halpern’s semantics of causal claims can be easily adapted to account for the data we consider; namely, by dropping his ‘minimality’ condition (§4).

1 Introduction

Monotonicity offers an insightful window into the logical properties of natural language expressions. This is especially true of causal expressions. Taking entailment as the relevant order, two-place functions (such as determiners, and, in the case of causation, binary relations) can be investigated, in the terminology of Barwise and Cooper (1981), in terms of downward and upward monotonicity in their left and right arguments.

In this paper we investigate whether actual causal claims are downward monotone in their cause argument (DMC). That is, we study whether the truth of a causal claim is preserved under strengthening the cause, where strength is understood as logical entailment. The answer to this question is not immediately obvious. On the one hand, there are apparent counterexamples; for example, it is not at all clear whether (1a) entails (1b):

* I am grateful to my supervisor Katrin Schulz for her guidance, to Milica Denić and two anonymous reviewers for remarkably helpful comments, and to Simon Rey, with whom I first discussed the topic of this paper. Special thanks also go to the organizers of the TLLM workshop for persevering despite the coronavirus pandemic.

- (1) a. Alice flicking the switch caused the light to turn on.
b. # Alice flicking the switch and it raining in New Zealand caused the light to turn on.

A first guess why (1b) is unacceptable could be that the sentence is false, which would result, for example, if the semantics of *cause* does not allow causes to be stronger than strictly required for the causal claim to be true. However, sometimes the cause is stronger than required but the causal claim is still acceptable:

- (2) Reyna was born at Royal Bolton Hospital but received a Danish passport because her mother was born in Copenhagen.¹

Having a mother born in Copenhagen is not necessary for one to acquire a Danish passport. When it comes to receiving a Danish passport, there is nothing special about Copenhagen compared to anywhere else in Denmark.

In this paper we propose that causes are never too strong. In other words, causal claims are downward monotonic in their cause argument. Thus (1a) entails (1b), but this is not a counterexample to DMC because, while (1b) is true whenever (1a) is, in such cases (1b) is unassertable because it triggers the scalar implicature that (1a) is false (as argued for in section 3.2 below).

In this paper we concentrate on English causal claims, where we understand “causal claims” to be either of the form “*C* caused *E*” or “*E* because of *C*”. In what follows we will consider both constructions, putting aside some evidence that there might be subtle differences in meaning between them.²

It is worth investigating the monotonicity properties of causal claims for two reasons. The first is that while there is a great deal of research on the monotonicity properties of quantifiers (beginning with the influential work of Barwise and Cooper 1981, van Benthem 1984 and Keenan and Stavi 1986), comparatively little has been written about the monotonicity properties of natural language connectives. It might be objected that the monotonicity properties of connectives are so straightforward that there is nothing much to say (e.g. clearly negation is downward entailing, and conjunction and disjunction are upward monotone in their left and right arguments). However, the connective *because* presents a particularly complex case study to test whether generalizations claimed to hold for determiners—e.g. that all simple determiners are monotone (Barwise and Cooper 1981)—also hold for connectives.

¹ *The Bolton News*, 12 February 2020. <https://www.theboltonnews.co.uk/news/18226923.bolton-born-woman-receives-british-passport-six-year-fight/>

² Copley and Wolff (2014: 55) offer the following example.

- (i) a. Lance Armstrong won seven Tours de France because of drugs.
b. Drugs caused Lance Armstrong to win seven Tours de France.

According to Copley and Wolff, (ia) is true but (ib) is false. We will not attempt to theorize any difference in meaning between (ia) and (ib) here, and will consider both constructions with *cause* and with *because* below.

The second reason to investigate the monotonicity properties of causal claims is that they can teach us about the semantics of causal claims more generally. Any semantics of causal claims should be able to say something about problematic cases such as (1b) where the cause is stronger than required for the claim to be true. Resolving the status of such sentences is important for the semantics of causal claims in general.

The structure of the paper is as follows. In section 2 we present data for and against DMC in causal claims. Section 3 shows that the data is readily accounted for in terms of the pragmatics of causal claims; in particular, by attending to their scalar implicatures. In section 4 we investigate DMC in the semantics of actual causal claims proposed by Halpern (2016). We show that the validity of DMC depends on how Halpern structures the variables in his modeling framework, that of structural causal models (Pearl 2000). We end by showing that there is reason for Halpern to modify his framework to support the proposal that causal claims are always DMC, by dropping a condition he calls ‘minimality’.

Before we proceed, we must define exactly what it means for one causal claim to entail another. This might seem straightforward, but the task is complicated by the presuppositions of causal claims. Let us turn to those presuppositions now.

1.1 Taking the soft presuppositions of causal claims into account

Causal claims appear to presuppose that their propositional arguments are true. For example, the sentences in (3) presuppose that the mentioned causes and effects actually occurred (e.g. that Joe Kennedy advanced, had legal skills and that his bosses were starstruck).

- (3)
- a. Did Joe Kennedy advance because of his legal skills or because his bosses were starstruck?³
 - b. The parents of Oscar Knox have said their son didn’t die because he had cancer but because they ran out of options to treat it.⁴
 - c. Did hospital readmissions fall because per capita admission rates fell?⁵

However, causal claims are still felicitous when the common ground does not establish that the stated cause or effect occurred, as shown in (4). For instance, (4b) does not imply that Putin had a stroke, and (4d) does not imply that the death rate dropped in Chicago.

³ *Boston Magazine*, 13 May 2020. <https://www.bostonmagazine.com/news/2020/05/13/joe-kennedy-iii-profile/>

⁴ *Irish News*, 9 September 2017. <https://www.irishnews.com/news/2017/09/09/news/family-of-oscar-knox-establish-charity-in-son-s-memory-1132115/>

⁵ *Health Affairs*, November 2019. <https://www.healthaffairs.org/doi/abs/10.1377/hlthaff.2019.00411>

- (4)
- a. The outcry which followed *Morgan* was not because the House of Lords had changed the law but because the public mistakenly thought it had done so.⁶
 - b. Did a stroke cause Putin’s awkward English?⁷
 - c. If a mechanical failure caused my injury, can I still sue?⁸
 - d. No, the coronavirus did not cause the death rate to drop in Chicago... Overall, deaths don’t appear to be declining.⁹
 - e. Did NJ bail reform cause a surge in crime? ... Concerns about a possible spike in crime did not materialize.¹⁰
 - f. Dogs do not have ears because they have anything we don’t. They have ears because they have ears.¹¹ (Linebarger 1987)

The data in (4) suggest that causal claims ‘softly’ presuppose in the sense of Abusch (2002, 2010) that their propositional arguments are true, where soft triggers are “presupposition triggers where the presuppositional behavior is weak and easily suspendable” (Abusch 2002). Romoli (2011, 2015) proposes in particular that *because* softly presupposes that its propositional arguments are true. Moreover, many authors have concluded that soft presuppositions are pragmatically derived (e.g. Simons 2001, Abusch 2002, 2010, Abbott 2006, Chemla 2009, Romoli 2015). For example, Abrusán (2016) explains the ‘soft–hard’ distinction using general principles governing the interaction of information structure and context.

While soft presuppositions are pragmatically derived, monotonicity properties are traditionally understood as part of an expression’s literal meaning, independent of pragmatic reasoning. For example, we say *every* is downward monotone in its restrictor: *every P is Q* implies *every P’ is Q* whenever $P' \subseteq P$. This is despite the fact that from an utterance of *every P’ is Q*, one would typically infer that *some P’ is Q*. (5a) entails (5b), even though there are contexts where (5a) is assertable but (5b) is not (e.g. when no students are over 70).

- (5)
- a. Every student passed the test.
 - b. \Rightarrow Every student over the age of 70 passed the test.

⁶ Jennifer Temkin, *Rape and the Legal Process*. Oxford University Press, 2002.

⁷ *The Atlantic*, 12 June 2013. <https://www.theatlantic.com/international/archive/2013/06/did-a-stroke-cause-putins-awkward-english/276824/>

⁸ <https://galliganlaw.com/2018/08/29/mechanical-failure-caused-injury/>

⁹ <https://www.politifact.com/factchecks/2020/apr/03/facebook-posts/no-coronavirus-did-not-cause-death-rate-drop-chica/>

¹⁰ <https://eu.northjersey.com/story/news/new-jersey/2019/04/02/nj-bail-reform-no-crime-surge-pretrial-release/3336423002/>

¹¹ (4f) shows that the presupposition of *because* can be suspended in a more subtle way than the other examples in (4). Chierchia (2013: 378) argues that the negative polarity item *any* in (4f) is acceptable in contexts where the presupposition/implicature of *because*—that dogs have something we don’t—does not arise. If it did, *any* would find itself in a non-downward entailing context and would therefore not be licensed according to Chierchia’s theory and the Fauconnier–Ladusaw hypothesis.

In defining the monotonicity properties of causal claims, we will take into account the inference that their propositional arguments are true. The definition of monotonicity properties for causal claims we adopt in this paper is given below.

Definition 1 (Downward monotonicity in the cause (DMC)). *We define that cause (respectively, because) is downward monotone in its cause if and only if the following holds for any propositions C , C^+ and E such that C^+ entails C .*

If C cause E (respectively, E because C) is true and C^+ is true, then C^+ cause E (respectively, E because C^+) is also true.

Since the inference that C^+ is true is likely pragmatically derived, this perspective represents a departure from how the monotonicity properties of natural language expressions are traditionally understood.¹² However, the move it is necessary to avoid trivializing the question whether *cause* and *because* are downward monotone in their causes. Triviality would result because without the underlined clause in Definition 1, we could find counterexamples to downward monotonicity simply by picking a false C^+ . For instance, the entailment from (1a) to (1b) (repeated below) would fail simply because there are cases where it is not raining in New Zealand.

- (1) a. Alice flicking the switch caused the light to turn on.
- b. Alice flicking the switch and it raining in New Zealand caused the light to turn on.

2 Data on DMC in causal claims

If causal claims are not downward monotone in their cause, it is because in some cases, the truth of a causal claim is not preserved under strengthening the cause. That is, DMC fails just in case there are causal claims where the cause is too strong for the causal claim to be true.

Are there cases where the cause is stronger than required for the claim to be true, but the causal claim is still assertable? We already saw an example of such an assertion in (2), repeated as (6a) below with further examples:

- (6) a. Reyna was born at Royal Bolton Hospital but received a Danish passport because her mother was born in Copenhagen. = (2)

¹² Note that the underlined clause in Definition 1 would not result from redefining monotonicity using Strawson entailment; that is, by redefining *cause* to be downward monotone in its cause iff C cause E Strawson entails C^+ cause E whenever C^+ entails C (where p Strawson entails q just in case whenever p is true and q is defined, q is true; see von Stechow 1999: 104). This is because C^+ cause E can be defined even when C^+ is false; e.g. given (4a), *The outcry was because the House of Lords had changed the law* is false—hence defined—even though the law did not in fact change. Thanks to Milica Denić for raising the issue of Strawson entailment.

- b. He has an American passport because he was born in Boston.¹³
- c. Naama Issachar ... could spend up to seven-and-a-half years in a Russian prison because 9.5 grams of cannabis were found in her possession during a routine security check.¹⁴
- d. A 90-day study in 8 adults found that supplementing a standard diet with 1.3 cups (100 grams) of fresh coconut daily caused significant weight loss.¹⁵

For example, (6c) is acceptable even though presumably, Naama Issachar would still have gone to prison if she had been caught with, say, 9 grams of cannabis.

To take a more extreme example, the causes in (7) are far stronger (in the sense of logical entailment) than required to make the effect occur, yet the causal claims are still assertable.

- (7) a. Computers do an awful lot of deliberation, and yet their every decision is wholly caused by the state of the universe plus the laws of nature.¹⁶
- b. If anything is happening at this moment in time, it is completely dependent on, or caused by, the state of the universe, as the most complete description, at the previous moment.¹⁷
- c. If you keep asking “why” questions about what happens in the universe, you ultimately reach the answer “because of the state of the universe and the laws of nature.”¹⁸

If causal claims were not DMC, it would mean there are contexts where C cause E is true but C^+ cause E is false for some C and C^+ where C^+ entails C . In other words, we would expect some true causal claim to become false by making the cause too strong. Though in (7) we find causal claims where C^+ is as strong as it can possibly be, but the claim is still assertable. Assuming that the speakers are following Grice’s maxim of quality (Grice 1975), the speakers of these sentences take them to not only be assertable, but also true.

Now, the sentences in (6) and (7) do not provide conclusive evidence that causal claims are DMC. One could reply that we have missed the cases where a true causal claim is made false by strengthening the cause. Nonetheless, the data in (6) and (7) pose a challenge: one who believes that some causal claims are made false by strengthening the cause, and seeks to explain why, must ensure that their explanation does not also predict the falsity of the examples above.

¹³ https://rupaulsdragrace.fandom.com/wiki/Charlie_Hides

¹⁴ *The Jerusalem Post*, 24 November 2019. <https://www.jpost.com/israel-news/will-putin-release-issachar-before-he-visits-israel-in-january-analysis-608884>

¹⁵ <https://www.healthline.com/nutrition/coconut-meat>

¹⁶ <http://commonsenseatheism.com/?p=899>

¹⁷ George Ortega, *Exploring the Illusion of Free Will*, 2013. <http://causalconsciousness.com/Second%20Edition%20Chapters/14.%20%20Why%20Both%20Causality%20and%20Randomness%20Make%20Free%20Will%20Impossible.htm>

¹⁸ <https://www.edge.org/response-detail/10164>

3 Explaining apparent failures of DMC

3.1 A possible explanation of the failure of DMC

In (1) we saw initial evidence that actual causal claims are not always downward monotone in their cause arguments. Let us consider again the contrast observed in (1), repeated below.

- (1) a. Alice flicking the switch caused the light to turn on.
- b. Alice flicking the switch and it raining in New Zealand caused the light to turn on.

If causal claims are indeed not DMC, one might seek to explain this property in terms of counterfactual dependence. Beginning with Hume (1748: section VII) and taken up again by Lewis (1973), counterfactual dependence analyses of causation seek to analyse causal claims in terms of the counterfactual, *if the cause had not occurred, the effect would not have occurred* (though this view is plagued by a host of counterexamples, see e.g. Paul 1998, Schaffer 2000, Hall and Paul 2003: and many more).

In much recent work on counterfactuals, counterfactual antecedents can raise multiple scenarios, and a counterfactual is true just in case the consequent holds in *every* scenario raised by the antecedent (Kratzer 1986, Alonso-Ovalle 2006, von Stechow 2001, Ciardelli 2016 as well as many others, though see Stalnaker 1981 for an alternative view). Under this assumption, counterfactual dependence analyses of the semantics of causal claims make the following prediction:

- (8) (1b) is true iff in **all** scenarios raised by the antecedent $\neg(\text{Alice flick switch} \wedge \text{rain in NZ})$, the light turns on.

With this apparatus, one could explain that (1b) is unassertable because it is false, and that it is false because, if it had not been that Alice flicked the switch and it was raining in New Zealand, there are multiple scenarios to consider. In particular, in one scenario raised by the antecedent, where it does not rain in New Zealand but Alice still flicks the switch, the light still turns on. (1b) would therefore be predicted to be false because the counterfactual dependence claim fails: in some scenario raised by the antecedent, *If the cause had not occurred*, the effect still occurs.

However, this explanation makes the wrong prediction for the sentences in section 2. It wrongly predicts the sentences in (6) and (7) to be false. For example, if Renya's mother hadn't been born in Copenhagen, Renya might have still received a Danish passport, say, if her mother had been born in Aarhus instead. And (7b) would be false because, taking anything that is happening at this moment in time (e.g. the bird flying outside my window), if the state of the universe at the previous moment had been different, there are many possibilities to consider. Presumably in some of these, the bird is still flying outside my window.

Since the above explanation in terms of counterfactual dependence cannot account for the fact that the sentences in (6) and (7) are assertable but (1b) is not,

let us examine an alternative approach. This account will attend to differences in the sentences' implicatures.

3.2 Pragmatic deviance via false implicature

While (1) purports to show that causal claims are not DMC, an alternative response is that (1b) is true but unassertable because it has a false implicature. Without appealing to DMC, one could seek to explain that (1b) falsely implicates the existence of a causal relationship between New Zealand's weather and the light. For, under standard assumptions about the calculation of alternatives (e.g. via deletion, see Katzir 2007), (1a) is a competing alternative utterance to (1b). So after an utterance of (1b), a listener would naturally attempt to construct a reason for mentioning the weather in New Zealand; for example, that there is in fact a causal relationship between the weather in weather and the light. The pragmatic deviance of (1b) makes it hard to conclude from examples like (1) that causal claims violate DMC. Indeed, against expectations, examples like (1) may even provide evidence that causal claims are DMC after all. We pursue this idea next.

The above pragmatic explanation of the unassertability of (1b) was admittedly rather vague. We did not provide a precise account of how (1b) implicates that the weather in New Zealand is 'causally relevant' to the light, nor what notion of 'causal relevance' is at work in the pragmatic calculation. Such an explanation could appeal to the maxim of relevance, though it is unclear how exactly the explanation would proceed. In contrast, if causal claims are DMC, it is easy to derive exactly why (1b) is unassertable: it has a false scalar implicature. Given that (1a) is an alternative utterance to (1b) (created by deleting material from (1b)), if causal claims are DMC then (1a) entails (1b), in which case a speaker who opts for (1b) is using a weaker utterance when a stronger alternative, (1a), is available. If a listener believes that a speaker of (1b) is obeying the maxim of quantity, the listener would infer that the speaker believes (1a) to be false.

Thus what turned out to be an apparent counterexample to DMC can actually be construed an argument in its favor. If *cause* is DMC, then $C \text{ cause } E$ entails $(C \wedge D) \text{ cause } E$. The explanation of (1b)'s unassertability thus becomes exactly parallel to the explanation why (9a) is unassertable when it is common ground that all students passed the test; namely, both sentences are literally true but have a false scalar implicature.

- (9)
 - a. Some students passed the test.
 - b. **Implicates:** Not all students passed the test.

- (10)
 - a. Alice flicking the switch and it raining in New Zealand caused the light to turn on. =(1b)
 - b. **Implicates:** \neg (Alice flicking the switch caused the light to turn on).

While it may be possible to derive the infelicity of (1b) without assuming that causal claims are DMC (for example, by appealing to the maxim of relevance) the assumption of DMC allows us to derive the infelicity of (1b) ‘out of the box’, so to speak, from the familiar mechanism of scalar implicature calculation.¹⁹

Now that we have a proposed explanation for the unassertability of examples like (1b), let us put that theory to the test. We do so in the following two sections.

3.3 Sensitivity to alternatives

An utterance’s pragmatically enriched meaning, unlike its at-issue contribution, is calculated by taking into account what the speaker could have said instead—the utterance’s *alternatives*. If causal claims where the cause is stronger than strictly required for the claim to hold such as (1b) are true, but unassertable due to a false scalar implicature, we would expect it to be assertable in contexts where the alternatives are such that no false implicature arises.

This prediction is borne out. We find evidence in the examples from section 2. Consider (2), repeated below (though note that the remarks in this section could apply equally well to any of the sentences in (6) or (7)):

- (2) Reyna was born at Royal Bolton Hospital but received a Danish passport because her mother was born in Copenhagen.

If *Denmark* were an alternative to *Copenhagen* in (2), then assuming DMC, we would expect (2) to trigger the scalar implicature that it is false that Reyna received a Danish passport because her mother was born in Denmark. This is because under DMC we have the entailment:

- (11) a. Reyna received a Danish passport because her mother was born in Denmark. *E because C*
 b. \Rightarrow Reyna received a Danish passport because her mother was born in Copenhagen. *E caused C⁺*

We can account for the assertability of (2) by assuming that *Denmark* is not an alternative to *Copenhagen* in (2) and therefore does not trigger a false implicature. To put this explanation to the test, we can alter the sentence to force

¹⁹ For this explanation to work, the scalar implicature calculation must be *obligatory* and *blind to contextual information* (in the sense of Magri 2009). The implicature must be obligatory because if it could be canceled—say, because the truth of (1a) is already in the common ground, which is inconsistent with the implicature—we would expect (1b) to be assertable, contrary to observation (assuming (1b) is not unassertable for some other reason). And the implicature calculation must be blind to contextual information for the following reason. Assuming (1a) is in the common ground, then by DMC, (1b) is too. So (1a) and (1b) are contextually equivalent—true in all the same worlds compatible with the common ground. But then if scalar implicatures were calculated with respect to contextual entailment, (1a) would not be a strictly more informative alternative to (1b), no false implicature would be generated, and we would instead expect (1b) to be assertable (again, assuming (1b) is not unassertable for some other reason).

Denmark to be an alternative to *Copenhagen* and check whether the scalar implicature is triggered as predicted. Following the theory of alternative calculation from Fox and Katzir (2011), we can make *Denmark* an alternative by making it contextually salient and focusing *Copenhagen*, as in the following dialogue, where subscript F indicates focus marking:

- (12) a. A: I have a Danish passport because my father was born in Denmark. Why do you have one?
 b. B: ??Because my mother was born in [Copenhagen]_F.

In this context, (12b) indeed triggers the implicature that *Copenhagen* is somehow special when it comes to receiving Danish passports; in other words, that it is not true that B has a Danish passport because their mother was born in Denmark. This is correctly predicted by the entailment in (11), an entailment guaranteed by DMC.

Note that while (2) optionally triggers a false scalar implicature, (1b) does so obligatorily:

- (1b) # Alice flicking the switch and it raining in New Zealand caused the light to turn on.

As we saw in 3.2, we can account for this by assuming DMC and that *Alice flicking the switch* is obligatorily an alternative to *Alice flicking the switch and it raining in New Zealand*.²⁰

Thus DMC allows us to explain why the sentences in (6) and (7) are assertable while (1b) is not. The difference lies in how their alternatives are derived. (6) and (7) are assertable provided that no weaker cause is an alternative to the cause appearing in the sentence, in which case no false implicature is triggered, while (1b) is obligatorily unassertable when (1a) is true because *C* is obligatorily an alternative to *C ∧ D* (e.g. via deletion; see Katzir 2007), meaning *C ∧ D cause E* obligatorily triggers the scalar implicature $\neg(D \text{ cause } E)$, that (1a) is false.

3.4 Behavior in downward entailing environments

One of the most straightforward ways to test whether sentence (1b) is false, or true but unassertable, is to put it in a downward entailing environment. Examples are shown in (13).

- (13) a. ?? I doubt that the light turned on because Alice flicked the switch and it was raining in New Zealand.
 b. ?? No one thinks that Alice flicking the switch and it raining in New Zealand caused the light to turn on.

In this subsection we argue that sentences in (12) provide evidence against the hypothesis that the embedded causal claim (1b) is false, and in favor of the hypothesis that (1b) is true but unassertable due to a false scalar implicature.

²⁰ For further discussion of the obligatory nature of this implicature, see footnote 19.

The crucial observation is that the sentences in (13) are improved with prosodic focus on *and it (was) raining in New Zealand*. This is unexpected according to a theory where (1b) is literally false, and so the sentences in (13) should be straightforwardly true. However, this is expected if (1b) is false but can be rescued by metalinguistic negation targeting a scalar implicature triggered by the focused material. We develop this proposal below.

Examples of metalinguistic negation are shown in (14):

- (14) a. He didn't eat [some]_F of the cookies. He ate [all]_F of them.
 b. I don't [like]_F scallops. I [love]_F them.

Metalinguistic negation is used to target an utterance's non-at-issue content. In (14), metalinguistic negation targets the scalar implicatures triggered by the focused material, with *some* implicating *not all* and *like* implicating *don't love*.

Let us consider some more clear-cut examples of metalinguistic negation in causal claims. In (14) and (15) alike, the focus marking is obligatory for the sentences to be felicitous.

- (15) a. I refuse to eat it, not because it's a [pineapple]_F pizza, but because it's [pizza]_F. I hate pizza.
 b. I am not upset because you lost my wedding ring [and my phone]_F. I'm upset because you lost [my wedding ring]_F.
 c. The fact that the meeting [happened]_F caused my surprise. It wasn't the fact that the meeting happened [on a Sunday]_F.

According to Horn (1985, 1989) and Burton-Roberts (1989), metalinguistic negation only applies after the hearer realizes the sentence cannot be interpreted using truth-functional, descriptive negation. A straightforward explanation why descriptive negation cannot apply in (14) and (15) is that the negated claim is entailed by the clause following it. For if the entailment relations in (16) hold, applying descriptive negation to the stronger claim would result in a contradictory meaning.

- (16) a. He ate all of the cookies. \Rightarrow He ate some of the cookies.
 b. I love scallops. \Rightarrow I like scallops.

Similarly, assuming DMC the following entailments hold:

- (17) a. The light turned on because Alice flicked the switch. \Rightarrow The light turned on because Alice flicked the switch and it was raining in New Zealand.
 b. I am upset because you lost my wedding ring. \Rightarrow I am upset because you lost my wedding ring and my phone.
 c. I refuse to eat it because it's pizza. \Rightarrow I refuse to eat it because it's pineapple pizza.
 d. The fact that the meeting happened caused my surprise. \Rightarrow The fact that the meeting happened on Sunday caused my surprise.

An alternative perspective on metalinguistic negation proposes that there is only one kind of negation, but it can target an utterance’s pragmatically enriched meaning (Carston 1996, 2002, Noh 1998, 2000, Moeschler 2019). If causal claims are DMC—and so the entailment relations in (17) hold—one can apply the scalar implicature calculation proposed in section 3.2 to predict the following implicatures.

- (18) a. The light turned on because Alice flicked the switch and it was raining in New Zealand.
 Scalar implicature: \neg (The light turned on because Alice flicked the switch.)
- b. I refuse to eat it because it’s pineapple pizza.
 Scalar implicature: \neg (I refuse to eat it because it’s pizza)
- c. I am upset because you lost my wedding ring and my phone.
 Scalar implicature: \neg (I am upset because you lost my wedding ring)
- d. The fact that the meeting happened on Sunday caused my surprise.
 Scalar implicature: \neg (The fact that the meeting happened did not cause my surprise)

Adopting the theory of metalinguistic negation of Carston (1996, 2002), Noh (1998, 2000), Moeschler (2019), we can explain the data in (15) as a case where the negation targets the causal claims’ scalar implicatures.

Thus, regardless of which perspective on metalinguistic negation we take, we are able to explain the observation that the sentences in (13) and (15) require focus marking to be felicitous, following the pattern of more familiar examples of metalinguistic negation such as (14). Crucially, this explanation requires assuming that the entailment relations in (17) hold—a consequence of DMC. The fact that (13) and (15) pattern with other examples of metalinguistic negation therefore provides further support for DMC.

4 Truth conditions for causal claims: Halpern (2016)

The data in section 2 provided evidence that causal claims are DMC. In this section we show that a recent influential analysis of the truth conditions of causal claims, due to Halpern (2016), does not account for this fact. However, we show that Halpern’s semantics of causal claims can be easily adapted to account for the data we consider; namely, by dropping his ‘minimality’ condition.

4.1 Halpern’s semantics for causal claims

Halpern (2016) defines his truth conditions for causal claims in terms of structural causal models (Pearl 2000).²¹ Let us briefly review this framework. We let

²¹ A reviewer rightly asks how causal network models fit with natural language semantics, and in particular how the network is supposed to be derived from natural language utterances (e.g. Does the network come from explicit text? From implicit

V be a set of variables of arbitrary arity, and where X is a variable, let $\mathcal{R}(X)$ denote the *range* of X , that is, the set of values X may take. A structural causal model is then defined as follows.

Definition 2 (Structural causal model). *A structural causal model is a triple $M = (V, E, F)$ where V is a set of variables, (V, E) is a directed acyclic graph, and F is a set of functions of the form*

$$F_X : \mathcal{R}(pa_X) \rightarrow \mathcal{R}(X),$$

one for each endogenous variable $X \in V$ (X is endogenous iff X has a parent in the graph), where $pa_X := \{Y \in V \mid (Y, X) \in E\}$ is the set of parents of X in the graph (V, E) .

Where \bar{U} is the set of exogenous (i.e. parentless) variables in (V, E) , and $\vec{u} \in \mathcal{R}(\bar{U})$, we call \vec{u} a setting of M .

In the structural causal modeling framework, the semantics of causal claims is understood in terms of interventions. An intervention is an operation that sets the value of a variable X by manually changing its function F_X . This is given in Definition 3.

Definition 3 (Truth conditions for interventions). *Where $M = (V, E, F)$ is a structural causal model, $M_{\vec{X}=\vec{x}}$ is the model (V, E, F') that results by setting, every variable $Y \in \vec{X}$, $F'_Y(\vec{z}) = y$ for every value \vec{z} of the parents of y .*

We write $(M, u) \models \vec{X} = \vec{x}$ just in case \vec{X} has value \vec{x} according to the equations in F , and write

$$(M, \vec{u}) \models [\vec{X} \leftarrow \vec{x}]Y = y \quad \text{iff} \quad (M_{\vec{X}=\vec{x}}, \vec{u}) \models Y = y.$$

With a treatment of interventions at hand, Halpern proposes the following truth conditions for causal claims.²²

Definition 4 (Halpern’s truth conditions for actual causal claims). *Let $M = (V, E, F)$ be a structural causal model, \vec{u} a context for M , and \vec{X} a vector of variables. $\vec{X} = \vec{x}$ is an actual cause of φ in the causal setting (M, \vec{u}) iff*

context?). In section 4.2 we will address one issue affecting the construction of the network: the choice of variables; in particular, how fine-grained we should take the variables to be. Unfortunately a larger assessment of the adequacy of causal networks in natural language semantics is beyond the scope of this paper. Though since much recent work in natural language semantics adopts causal networks as a model—especially in the semantics of conditionals (e.g. Schulz 2011, Briggs 2012, Ciardelli et al. 2018, Santorio 2019)—the question of their adequacy in natural language semantics arises for a number of authors.

²² Halpern actually proposes three separate versions of AC2: an ‘original’ definition, an ‘updated’, and a ‘modified’ definition. The modified version is what appears above. Halpern acknowledges that the original version is subject to counterexamples (Halpern 2016: example 2.8.1), and states that his “current preference” is for the modified definition. For this reason we only consider the modified definition.

AC1 $(M, \vec{u}) \models \vec{X} = \vec{x}$ and $(M, \vec{u}) \models \varphi$.

AC2 There is a vector \vec{W} of variables and a value \vec{x}' of \vec{X} such that

$$(M, \vec{u}) \models \vec{W} = \vec{w} \quad \text{and} \quad (M, \vec{u}) \models [\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}] \neg \varphi.^{23}$$

AC3 \vec{X} is minimal; there is no strict subset \vec{X}' of \vec{X} such that $\vec{X}' = \vec{x}'$ satisfies conditions AC1 and AC2, where \vec{x}' is the restriction of \vec{x} to the variables in \vec{X}' .

In essence, the three conditions state the following.

1. The cause and the effect actually occurred.
2. Fixing some variables to their actual values, if the cause had had a different value, the effect would not have occurred.
3. If the cause were any weaker (in the sense of logical entailment) it would not satisfy AC2.

While Halpern’s definition is phrased in terms of $\vec{X} = \vec{x}$ being “an actual cause” of φ , we will apply his analysis to the constructions considered in this paper: the verb *cause* and the connective *because*. One reason why it is worth examining how Halpern’s analysis fares with such constructions is that they occur much more frequently than either *a cause of* or *the cause of*.²⁴

4.2 An obstacle in the way of representing monotonicity in structural causal models

There is one theory-internal obstacle getting in the way using structural causal models to test the monotonicity of properties of causal claims. The problem is

²³ Strictly speaking, the condition AC2 above is not the condition proposed by Halpern (2016: 25). The condition above uses a conjunction, whereas Halpern’s own condition uses a conditional, requiring that there is a set \vec{W} of variables and a setting \vec{x}' of the variables in \vec{X} such that *if* $(M, \vec{u}) \models \vec{W} = \vec{w}$ *then* $(M, \vec{u}) \models [\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}] \neg \varphi$. The problem with the *if-then* formulation is that it predicts AC2 to always be true. Halpern’s formulation of AC2 is true whenever the antecedent is false, that is, whenever there is a set of variables and an assignment that is false in the actual context \vec{u} . But the actual context \vec{u} always makes some assignment of values to variables false, so Halpern predicts AC2 to be always true. I think Halpern simply miswrote the formula, and intended to write AC2 with a conjunction instead. I have therefore taken the liberty to rewrite his definition as it appears above.

²⁴ Searches of the British National Corpus (BNC) and Corpus of Contemporary American English (CCAE) reveal that for every occurrence of either *a cause* or *the cause* there are approximately 3 occurrences of *caused* (in both the BNC and CCAE) and 36 (BNC) and 62 (CCAE) occurrences, respectively, of *because*. Frequency of *a cause*: 609 (BNC), 4852 (CCAE); *the cause*: 2161 (BNC), 16586 (CCAE); *caused*: 9243 (BNC), 62527 (CCAE); *because*: 99496 (BNC), 1346051 (CCAE). Corpora accessed at <https://www.english-corpora.org/bnc/> and <https://www.english-corpora.org/coca/> on 5 October 2020.

that the variables in structural causal models are taken to be *logically independent*, in the sense that every assignment of values to variables is consistent.²⁵ One reason for the assumption of logical independence is that SCMs are typically employed to represent the effects of interventions (see Pearl 2000). Logical independence in the sense above is required for the effect of every intervention on an SCM to be defined.²⁶ The assumption of logical independence implies that (19a) and (19b) cannot both be analyzed in the same SCM; for if they could, it would be possible to intervene to have John born in Boston but not in the United States, contradicting the fact that Boston is in the United States.

- (19) a. John has an American passport because he was born in the United States.
 b. John has an American passport because he was born in Boston.

There are many ways one might propose to get around the problem of contradictory interventions. One way would be to take variables to be maximally fine-grained. For example, instead of a binary variable representing *Was John born in Boston?* we could use a variable with a much more fine-grained range representing *Where was John born?* By packaging logically dependent values inside the same variable, one avoids the problem of contradictory interventions because one cannot intervene to set the same variable to two different values.

Taking the variables to be fine-grained is one way to solve the problem of contradictory interventions. Though if we adopt fine-grained variables, we must make a slight technical modification to Definition 4 to adequately represent the sentences discussed in 2 in Halpern’s framework. The reason is that Halpern’s definition takes a cause to be an assignment of a *single* value to a variable (or vector of variables). Even if the variables are maximally specific, our ordinary causal talk often is not. The solution is straightforward enough: allow causes in Halpern’s definition to be sets of values, rather than a single value. For instance, if X represents where John was born, we might take $\mathcal{R}(X)$ to be a set of locations and let *Boston* and *United States* be the appropriate subsets of $\mathcal{R}(X)$. We can then express causes of varying specificity, for example, that $X \in \textit{Boston}$ caused John to have a US passport, or that $X \in \textit{United States}$ did. The changes to Definition 4 are given below.

Definition 5 (Allowing weaker causes in Halpern’s framework). *Where \vec{X} is a vector of variables and $A \subseteq \mathcal{R}(\vec{X})$, we say $\vec{X} \in A$ is an actual cause of φ in the causal setting (M, \vec{u}) iff*

- AC1’** $(M, \vec{u}) \models \vec{X} = \vec{x}$ for some $\vec{x} \in A$ and $(M, \vec{u}) \models \varphi$.
AC2’ There is a set \vec{W} of variables and a setting \vec{x}' of the variables in \vec{X} such that $\vec{x}' \notin A$, $(M, \vec{u}) \models \vec{W} = \vec{w}$ and $(M, \vec{u}) \models [\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}] \neg \varphi$.

²⁵ By ‘consistency’ here we mean consistency with logic and with analytic relations given by world knowledge—e.g. that Copenhagen is in Denmark—while allowing for inconsistency with the causal laws, represented by structural equations (Pearl 2000).

²⁶ Though see Beckers and Halpern (2018) for a proposal to restrict interventions to ‘allowable interventions’.

AC3' No subset \vec{X}' of \vec{X} also satisfies AC1' and AC2'.

According to Definition 5, causal claims are DMC with respect to causes that share the same variables. More exactly, we have the following, which is a straightforward consequence of the fact that if $A^+ \subseteq A$ and $\vec{x}' \notin A$, then $\vec{x}' \notin A^+$.

Fact 1. *For any causal model M and setting \vec{u} , according to Definition 5, if $\vec{X} \in A$ is an actual cause of φ in (M, \vec{u}) and $A^+ \subseteq A$, then $\vec{X} \in A^+$ is also an actual cause of φ in (M, \vec{u}) .*

By Fact 1, (19a) entails (19b), provided that *John was born in the United States* is represented by the same variable as *John was born in Boston*.

- (19) a. John has an American passport because he was born in the United States.
 b. \Rightarrow John has an American passport because he was born in Boston.

4.3 Failures of DMC in Halpern's framework: minimality

While Halpern predicts that causal claims are DMC for causes that are represented by the same variables, it turns out the opposite holds for the causes that are not represented by the same variables.

Fact 2. *For any causal model M and setting \vec{u} , according to Definition 5, if $\vec{X} \in A$ is an actual cause of φ in (M, \vec{u}) and $\vec{X} \subsetneq \vec{Y}$, then for no $B \subseteq \mathcal{R}(\vec{Y})$ is $\vec{Y} \in B$ an actual cause of φ in (M, \vec{u}) .*

Fact 2 holds because of Halpern's minimality condition. If \vec{X} and \vec{Y} were both actual causes of φ and $\vec{X} \subsetneq \vec{Y}$, then \vec{Y} would violate minimality (AC3'). Indeed, Halpern states that he added his minimality condition precisely to rule out such cases.

AC3 is a minimality condition, which ensures that only those elements of the conjunction $\vec{X} = \vec{x}$ that are essential are considered part of a cause; inessential elements are pruned. Without AC3, if dropping a lit match qualified as a cause of the forest fire, then dropping a match and sneezing would also pass the tests of AC1 and AC2. AC3 serves here to strip "sneezing" and other irrelevant, over-specific details from the cause. (Halpern 2016: 23)

Halpern's theory predicts that such "irrelevant, over-specific details" only make a truth conditional difference when they are represented by a separate variable. Overly specific causes do not render a causal claim false, provided the overly specific detail is still represented by the same variable as a weaker cause satisfying AC1–3. There is reason to think, however, that minimality should not be part of the truth conditions of causal claims after all. We explore a counterexample to minimality in the next section.

4.4 Against minimality

If we take Halpern’s definition of actual causality as an analysis of the verb *cause* or the connective *because*, his minimality condition leads to some surprising results. Consider the following scenario.²⁷ A committee is tasked with approving new company policies. The committee has two members: the Chairperson and the CEO. A policy is approved just in case both committee members approve it. Recently, a new proposal came before the committee. Independently, the Chairperson and CEO each liked the proposal, and so each voted in favor of adopting it.

- (20) a. The fact that the Chairperson voted ‘Yes’ and CEO voted ‘Yes’ caused the proposal to pass.
 b. The proposal passed because the Chairperson voted ‘Yes’ and the CEO voted ‘Yes’.




Fig. 1. A simple model of the voting scenario in (20)

We represent the sentences in (20) in Halpern’s framework as (21), ‘Agent = 1’ holds just in case the agent voted ‘Yes’, and ‘Result = 1’ holds just in case the policy was approved.

- (21) (Chair, CEO) = (1, 1) is an actual cause of Result = 1.

(21) clearly satisfies AC1. It satisfies AC2 because, taking W to be empty, there is another setting \vec{x}' of $\vec{X} = (\text{Chair}, \text{CEO})$ such that $(M, \vec{x}) \models [\vec{X} \leftarrow \vec{x}'] \text{Result} \neq 1$; indeed, any setting of (Chair, CEO) besides (1, 1)—namely, (1, 0), (0, 1) or (0, 0)—would suffice.

Nonetheless, (21) is false according to Halpern’s definition because it violates minimality (AC3). Taking $X' = \text{Chair}$ or $X' = \text{CEO}$, we have that $X' = 1$ also satisfies AC1 and AC2. This example does not seem to fit Halpern’s motivation for adopting minimality; namely, to strip “irrelevant, over-specific details from

²⁷ An anonymous reviewer points out that the following example is isomorphic to the conjunctive forest fire scenario considered by Halpern (2016: example 2.3.1). We find the following committee example slightly more natural than Halpern’s, in which a forest will not burn if struck by lightning or if a lit match is dropped, but will burn if both happen. Of course, since the two examples have the same causal structure, Halpern’s example could be used here without affecting the conclusions we draw in this section.

the cause” (Halpern 2016: 23). Since (21) seems perfectly acceptable, but violates the minimality condition (AC3), one might recommend abandoning the minimality condition altogether.

4.5 Partial causes to the rescue?

The previous section showed that, in virtue of minimality, Halpern makes the wrong prediction for conjunctive causes, predicting that the conjunction *The Chairperson voting ‘Yes’ and the CEO voting ‘Yes’* is not a cause of the motion passing, against intuitions. However, one might reply that we have simply mis-translated natural language into Halpern’s formal system.²⁸ In particular, one might argue that we have overlooked *partial causes*. Halpern (2016: 25) defines that whenever $\vec{X} = \vec{x}$ is a cause of φ in context (M, \vec{u}) , each conjunct of $\vec{X} = \vec{x}$ is *part* of a cause of φ in (M, \vec{u}) . Halpern then offers the following remarks on the relationship between his definition and natural language:

What we think of as causes in natural language correspond to parts of causes, especially with the modified HP definition [Definition 4 above]. Indeed, it may be better to use a term such as “complete cause” for what I have been calling cause and then reserve “cause” for what I have called “part of a cause”. (Halpern 2016: 25)

Under this formalization of natural language, Halpern predicts that the CEO voting ‘Yes’ and the Chairperson voting ‘Yes’ are each, on their own, complete causes of the motion passing. Besides the fact that this is a strange use of the word ‘complete’, the fact remains that Halpern makes the wrong predictions for the conjunction $(\text{CEO} = 1) \wedge (\text{Chair} = 1)$, classifying it as neither a complete nor partial cause of the motion passing.

Thus, even when we take into account Halpern’s suggestions above about how to formalize natural language in his framework, his definition of actual causality is still unsuitable as an analysis of the verb *cause* or the connective *because*. This is because his definition yields the wrong results for conjunctive causes, as in (20). It predicts the sentences in (20) to be false—regardless whether we interpret *caused* in (20a) as ‘partially caused’ or ‘completely caused’, and regardless whether we interpret *because* in (20b) as ‘partially because’ or ‘completely because’.

The example in section 4.4 therefore further supports dropping the minimality condition from Halpern’s definition of actual causality. We end by quickly proving that dropping minimality indeed has the desired effect, resulting in truth conditions for actual causal claims that are downward monotone in their cause argument.

²⁸ Thanks to an anonymous reviewer for encouraging me to include a discussion of partial causes.

4.6 Without minimality: DMC restored

Without AC3, Halpern predicts that causal claims are always downward monotone in their cause. We have already shown this in the case when one uses the same variables to represent the stronger and weaker cause (Fact 1). Below we show this in cases where the stronger cause is not represented by the same variables as the weaker cause.

Fact 3 (Downward monotonicity of $AC1 \wedge AC2$). *If $\vec{X} = \vec{x}$ satisfies AC1 and AC2 with respect to φ and (M, \vec{u}) , then for any variables \vec{Y} such that $(M, \vec{u}) \models \vec{Y} = \vec{y}$, the conjunction $\vec{X} = \vec{x} \wedge \vec{Y} = \vec{y}$ satisfies AC1 and AC2 with respect to φ and (M, \vec{u}) .*

Proof. AC1 follows from the assumption that $(M, \vec{u}) \models \vec{Y} = \vec{y}$. For if $\vec{X} = \vec{x}$ satisfies AC1 and $(M, \vec{u}) \models \vec{Y} = \vec{y}$, then $(M, \vec{u}) \models \vec{X} = \vec{x} \wedge \vec{Y} = \vec{y}$.

And if $\vec{X} = \vec{x}$ satisfies AC2, then there is a setting \vec{x}' of \vec{X} such that $(M, \vec{u}) \models [\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}] \neg \varphi$ for some set of variables \vec{W} such that $(M, \vec{u}) \models \vec{W} = \vec{w}$.

Let \vec{y}' be the value of \vec{Y} under the intervention setting \vec{X} to \vec{x}' and \vec{W} to \vec{w} . That is, let $\vec{y}' \in \mathcal{R}(\vec{Y})$ be such that $(M, \vec{u}) \models \vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w} \vec{Y} = \vec{y}'$. Now, all structural causal models validate the following principle (which Pearl calls ‘consistency’, see Pearl 2000: Corollary 7.3.2):

$$\text{if } (M, \vec{u}) \models \vec{A} = \vec{a} \wedge \vec{B} = \vec{b} \text{ then } (M, \vec{u}) \models [\vec{A} \leftarrow \vec{a}] \vec{B} = \vec{b}.$$

Consistency says that intervening to set a variable to its actual value does not change the value of any variable. In particular, since $(M_{\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}}, \vec{u}) \models \vec{Y} = \vec{y}' \wedge \neg \varphi$, by consistency, $(M_{\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}}, \vec{u}) \models [\vec{Y} \leftarrow \vec{y}'] \neg \varphi$, which by Definition 3 is equivalent to

$$(M, \vec{u}) \models [(\vec{X}, \vec{Y}) \leftarrow (\vec{x}', \vec{y}'), \vec{W} \leftarrow \vec{w}] \neg \varphi.$$

Hence $(\vec{X}, \vec{Y}) = (\vec{x}, \vec{y})$ satisfies AC2 with respect to φ and (M, \vec{u}) .

Thus, without minimality, Halpern’s theory predicts that causal claims are always DMC.

5 Conclusion

While initial evidence suggests that causal claims are not DMC, the data can be explained by assuming that causal claims are in fact DMC. Assuming so allows us to explain the infelicity of the causal claims with a stronger cause as a case of false scalar implicature (section 3.2). We also saw though the phenomenon of metalinguistic negation in section 3.4 a parallel between paradigmatic entailments (e.g. *all* entails *some*, *love* entails *like*) and entailment relations between causal claims (*C* cause *E* entails *C*⁺ cause *E* whenever *C*⁺ entails *C*).

Turning to Halpern’s semantics of causal claims, we showed what whether causal claims are DMC according to Halpern (2016) depends on how the variables

are structured, though by making a slight modification to Halpern’s theory—abandoning minimality—Halpern predicts that causal claims are always DMC. The modification improves Halpern’s truth conditions for actual causal claims by allowing him to make the right predictions for claims with conjunctive causes (section 4.4), a benefit that cannot be achieved by interpreting the causal relation in question as either partial or complete (section 4.5).

While dropping minimality and validating DMC improves Halpern’s semantics of causal claims, the question remains whether the resulting theory yields a convincing formal theory of causation.²⁹ Recent work by Beckers and Vennekens (2018) suggests that there are more fundamental problems with Halpern’s analysis, problems that cannot be solved by dropping minimality. Nonetheless, while we have taken Halpern’s framework as an influential case study, the data presented above suggest that every semantics of causal claims should validate DMC. We leave it to future work to determine whether other analyses—such as Yablo (2002), Beckers and Vennekens (2018), Loew (2019), and Andreas and Günther (2020)—offer a satisfactory treatment of the monotonicity properties of causal claims.

Bibliography

- Barbara Abbott. Where have some of the presuppositions gone? *Drawing the boundaries of meaning: Neo-Gricean studies in pragmatics and semantics in honor of Laurence R. Horn*, 80:1–20, 2006.
- Márta Abrusán. Presupposition cancellation: explaining the ‘soft-hard’ trigger distinction. *Natural Language Semantics*, 24(2):165–202, 2016. <https://doi.org/10.1007/s11050-016-9122-7>.
- Dorit Abusch. Lexical alternatives as a source of pragmatic presuppositions. In *Semantics and linguistic theory*, volume 12, pages 1–19, 2002. <https://doi.org/10.3765/salt.v12i0.2867>.
- Dorit Abusch. Presupposition triggering from alternatives. *Journal of Semantics*, 27(1):37–80, 08 2010. <https://doi.org/10.1093/jos/ffp009>.
- Luis Alonso-Ovalle. *Disjunction in alternative semantics*. PhD thesis, University of Massachusetts Amherst, 2006. URL <http://people.linguistics.mcgill.ca/~luis.alonso-ovalle/papers/alonso-ovalle-diss.pdf>.
- Holger Andreas and Mario Günther. Causation in terms of production. *Philosophical Studies*, 177(6):1565–1591, 2020. <https://doi.org/10.1007/s11098-019-01275-3>.
- Jon Barwise and Robin Cooper. Generalized quantifiers and natural language. *Linguistics and Philosophy*, 4(2):159–219, Jun 1981. ISSN 1573-0549. <https://doi.org/10.1007/BF00350139>.
- Sander Beckers and Joseph Y. Halpern. Abstracting causal models. *CoRR*, abs/1812.03789, 2018. URL <http://arxiv.org/abs/1812.03789>.

²⁹ Thanks to an anonymous reviewer for raising this issue.

- Sander Beckers and Joost Vennekens. A principled approach to defining actual causation. *Synthese*, 195(2):835–862, Feb 2018. <https://doi.org/10.1007/s11229-016-1247-1>.
- Johan van Benthem. Questions about quantifiers 1. *The Journal of Symbolic Logic*, 49(2):443–466, 1984. <https://doi.org/10.2307/2274176>.
- Rachael Briggs. Interventionist counterfactuals. *Philosophical Studies*, 160(1): 139–166, Aug 2012. <https://doi.org/10.1007/s11098-012-9908-5>.
- Noel Burton-Roberts. On horn’s dilemma: presupposition and negation. *Journal of Linguistics*, 25(1):95–125, 1989. <https://doi.org/10.1017/S0022226700012111>.
- Robyn Carston. Metalinguistic negation and echoic use. *Journal of Pragmatics*, 25(3):309 – 330, 1996. [https://doi.org/https://doi.org/10.1016/0378-2166\(94\)00109-X](https://doi.org/https://doi.org/10.1016/0378-2166(94)00109-X).
- Robyn Carston. *Thoughts and utterances: The pragmatics of explicit communication*. John Wiley & Sons, 2002. <https://doi.org/10.1002/9780470754603>.
- Emmanuel Chemla. Similarity: Towards a unified account of scalar implicatures, free choice permission and presupposition projection. *Under revision for Semantics and Pragmatics*, 2009.
- Gennaro Chierchia. *Logic in grammar: Polarity, free choice, and intervention*. Oxford, 2013. <https://doi.org/10.1093/acprof:oso/9780199697977.001.0001>.
- Ivano Ciardelli. Lifting conditionals to inquisitive semantics. In *Semantics and Linguistic Theory*, volume 26, pages 732–752, 2016. <https://doi.org/10.3765/salt.v26i0.3811>.
- Ivano Ciardelli, Linmin Zhang, and Lucas Champollion. Two switches in the theory of counterfactuals. *Linguistics and Philosophy*, 2018. <https://doi.org/10.1007/s10988-018-9232-4>.
- Bridget Copley and Phillip Wolff. Theories of causation should inform linguistic theory and vice versa. In *Causation in Grammatical Structures*. Oxford University Press, 2014. <https://doi.org/10.1093/acprof:oso/9780199672073.003.0002>.
- Kai von Fintel. NPI licensing, Strawson entailment, and context dependency. *Journal of semantics*, 16(2):97–148, 1999. <https://doi.org/10.1093/jos/16.2.97>.
- Kai von Fintel. Counterfactuals in a dynamic context. In M. Kenstowicz, editor, *Ken Hale: A life in language*, volume 36, pages 123–152. MIT Press, 2001.
- Danny Fox and Roni Katzir. On the characterization of alternatives. *Natural language semantics*, 19(1):87–107, 2011. <https://doi.org/10.1007/s11050-010-9065-3>.
- Herbert Paul Grice. Logic and conversation. In *Speech acts*, pages 41–58. Brill, 1975.
- Ned Hall and Laurie A. Paul. Causation and preemption. *Philosophy of Science Today*, pages 100–130, 2003.
- Joseph Y Halpern. *Actual Causality*. MIT Press, 2016.
- Laurence R. Horn. Metalinguistic negation and pragmatic ambiguity. *Language*, 61(1):121–174, 1985. <https://doi.org/10.2307/413423>.

- Laurence R. Horn. *A Natural History of Negation*. Center for the Study of Language and Information, 1989.
- David Hume. *Philosophical Essays Concerning Human Understanding*. London: A. Millar, 1st edition, 1748.
- Roni Katzir. Structurally-defined alternatives. *Linguistics and Philosophy*, 30(6):669–690, Dec 2007. ISSN 1573-0549. <https://doi.org/10.1007/s10988-008-9029-y>.
- Edward L Keenan and Jonathan Stavi. A semantic characterization of natural language determiners. *Linguistics and philosophy*, 9(3):253–326, 1986. <https://doi.org/10.1007/BF00630273>.
- Angelika Kratzer. Conditionals. *Chicago Linguistics Society*, 22(2):1–15, 1986.
- David Lewis. Causation. *Journal of Philosophy*, 70(17):556–567, 1973. <https://doi.org/10.2307/2025310>.
- Marcia C Linebarger. Negative polarity and grammatical representation. *Linguistics and philosophy*, 10(3):325–387, 1987. <https://doi.org/10.1007/BF00584131>.
- Christian Loew. Causes as difference-makers for processes. *Philosophy and Phenomenological Research*, 98(1):89–106, 2019. <https://doi.org/10.1111/phpr.12424>.
- Giorgio Magri. A theory of individual-level predicates based on blind mandatory scalar implicatures. *Natural language semantics*, 17(3):245–297, 2009. <https://doi.org/10.1007/s11050-009-9042-x>.
- Jacques Moeschler. Representation and metarepresentation in negation. In K. Scott, B. Clark, and R. Carston, editors, *Relevance, Pragmatics and Interpretation*, pages 80–92. Cambridge, 2019. <https://doi.org/10.1017/9781108290593.008>.
- Eun-Ju Noh. Echo questions: Metarepresentation and pragmatic enrichment. *Linguistics and philosophy*, pages 603–628, 1998. <https://doi.org/10.1023/A:1005361528891>.
- Eun-Ju Noh. *Metarepresentation: A relevance-theory approach*, volume 69. John Benjamins Publishing, 2000. <https://doi.org/10.1075/pbns.69>.
- L. A. Paul. Keeping track of the time: Emending the counterfactual analysis of causation. *Analysis*, 58(3):191–198, 1998. <https://doi.org/10.1111/1467-8284.00121>.
- Judea Pearl. *Causality*. Cambridge University Press, 2000.
- Jacopo Romoli. The presuppositions of soft triggers aren't presuppositions. In *Semantics and linguistic theory*, volume 21, pages 236–256, 2011. <https://doi.org/10.3765/salt.v21i0.2619>.
- Jacopo Romoli. The Presuppositions of Soft Triggers are Obligatory Scalar Implicatures. *Journal of Semantics*, 32(2):173–219, 05 2015. <https://doi.org/10.1093/jos/fft017>.
- Paolo Santorio. Interventions in premise semantics. *Philosophers' Imprint*, 2019. URL <http://hdl.handle.net/2027/spo.3521354.0019.001>.
- Jonathan Schaffer. Trumping preemption. *Journal of Philosophy*, 97(4):165–181, 2000. <https://doi.org/10.2307/2678388>.

- Katrin Schulz. “If you’d wiggled A, then B would’ve changed”. *Synthese*, 179 (2):239–251, Mar 2011. <https://doi.org/10.1007/s11229-010-9780-9>.
- Mandy Simons. Disjunction and alternativeness. *Linguistics and Philosophy*, 24 (5):597–619, Oct 2001. <https://doi.org/10.1023/A:1017597811833>.
- Robert C. Stalnaker. *A Defense of Conditional Excluded Middle*, pages 87–104. Springer, 1981. https://doi.org/10.1007/978-94-009-9117-0_4.
- Stephen Yablo. De facto dependence. *The Journal of philosophy*, 99(3):130–148, 2002. <https://doi.org/10.2307/3655640>.