Illusions of Commutativity: The Case for Conditional Excluded Middle Revisited

Patrick Todd, Brian Rabern and Wolfgang Schwarz

1 Introduction

The principle of "Conditional Excluded Middle" (CEM) has been a matter of longstanding controversy in both semantics and metaphysics. The principle suggests (among other things) that for any coin that isn't flipped, there is a fact of the matter about how it *would* have landed if it *had* been flipped: either it would have landed heads, or it would have landed tails. This assumption has recently gained in popularity, as it appears to be supported by strong linguistic evidence. The evidence with which we will be concerned is that 'would' appears to commute with negation, meaning that 'not: if A, would C' appears to be equivalent to 'if A, would not C' – and the contention is that this evidence is best explained by the validity of CEM.

There is, however, a long list of operators that similarly appear to commute with negation, even though the corresponding excluded middle principles are indefensible. We will suggest that the data in support of CEM is best explained as a pragmatic effect.

We'll begin, in the next section, by reviewing the relevant data. In section 3, we will explain why many philosophers remain skeptical about CEM: the principle seems to have sweeping metaphysical implications. In sections 4 and 6, we will look at other operators that appear to commute with negation. Here, too, we will find philosophical reasons against taking appearances at face value. We will also find linguistic evidence against the commutativity assumption. The same kind of evidence, we argue, is available for conditionals. Finally, in section 7, we will consider how the data might be explained even if CEM isn't valid.

2 CEM: The data

We focus our attention on subjunctive conditionals.¹ In this context, Conditional Excluded Middle is the following schema, where 'A > C' expresses a subjunctive conditional with antecedent A and consequent C:

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

¹ Whether CEM holds for indicative conditionals is a separate matter about which we can remain neutral.

The question is whether all instances of this schema are valid. Most participants in this debate assume a principle of *conditional non-contradiction* CNC, at least for antecedents A that are deemed possible:²

(CNC)
$$\neg$$
((A > C) \land (A > \neg C)) if A is possible.

We will follow this assumption. Given this, CEM entails that 'would' commutes with negation (WCN):

(WCN)
$$\neg (A > C) \leftrightarrow (A > \neg C)$$
, if A is possible.

Conversely, the validity of (WCN) entails the validity of (CEM) whenever A is possible. Opponents of CEM hold that A>C and $A>\neg C$ can both be false, and not just if A is impossible. Thus opponents of CEM also reject WCN.

Informally, the WCN hypothesis says that denying a conditional A>C is equivalent to affirming the corresponding conditional A>¬C. This prediction is a little tricky to test because conditionals don't naturally occur under direct negation. We can, however, look at somewhat stilted sentences like (1a), which is arguably equivalent to (1b).

- (1) a. It is not the case that the coin would have landed heads if it had been flipped.
 - b. The coin would not have landed heads if it had been flipped.

A potentially better test uses verbs like 'doubt' or 'reject' that lexicalize negation. For example, one might intuit that (2a) and (2b) are equivalent ([Cariani and Santorio 2018]):

- (2) a. I doubt that the coin would have landed heads if it had been flipped.
 - b. I believe that the coin would not have landed heads if it had been flipped.

We can also look at negative answers to questions, as in (3).

- (3) If the coin had been flipped, would it have landed heads?
 - a. No.
 - b. No, it would not have landed heads.

The simple (3a) is intuitively equivalent to (3b).

Stronger evidence for WCN and CEM comes from quantified conditionals. [Williams 2010] considers (4a) and (4b).³

- (4) a. No student would have passed if they had goofed off.
 - b. Every student would have failed if they had goofed off.

² It is often assumed that subjunctive conditionals with impossible antecedents are (vacuously) true. On this account, A>C and A>¬C are both true if A is impossible. In this paper, we will not be interested in impossible antecedents.

^{3 [}Williams 2010: 652]. The argument builds on some observations of Higginbotham [Higginbotham 1986], and especially the discussion in [Von Fintel and Iatridou 2002]. Versions of this argument are likewise developed by [Mandelkern 2018] and [Cariani and Goldstein 2020]; it is similarly endorsed in [Goodman 2016].

Let's assume that "failing" and "passing" are the only possibilities for each student, so that 'failed' is equivalent to 'not passed'. In such a context, (4a) and (4b) are intuitively equivalent.

The relevant logical form of these sentences is plausibly something like (5a) and (5b).

```
(5) a. [No x: Sx](Gx > Px)b. [Every x: Sx](Gx > ¬Px)
```

Given that [No x: Sx] $\phi(x)$ is equivalent to [Every x: Sx] $\neg \phi(x)$, we get an equivalence between (6a) and (6b).

```
(6) a. [Every x: Sx]\neg(Gx > Px)
b. [Every x: Sx](Gx > \neg Px)
```

Trusting that the intuition of equivalence does not rest on specific features of the example, this suggests that any instance of (7a) is equivalent to the corresponding instance of (7b), at least if A is possible.

```
(7) a. [Every x: B]¬(A > C)b. [Every x: B](A > ¬C)
```

Notice that the only difference between these formulas concerns the position of the negation. The equivalence is expected if (WCN) and (CEM) are valid, and unexpected if they are invalid.⁴

In light of this data, Williams maintains that "denying CEM comes to seem rather heroic: one needs to explain away ... the apparent equivalences between quantified conditional statements." [Williams 2010: p.652]

We can be heroes. But before we set out on our heroic enterprise, we want to provide some philosophical motivation for doing so. Why might one be motivated to deny CEM?

3 Counterfacts?

The debate over CEM seems like an esoteric debate in the logic of conditionals. But the debate has important philosophical ramifications. Imagine 100 indistinguishable, fair, indeterministic coins that are never flipped. What would have happened if they had all been flipped? If WCN is valid then there is a unique and definite answer, something like: "The first would have landed tails, the second heads, the third heads, the fourth tails, ...". The true answer might not be this one, of course, but whatever it is, it similarly specifies the outcome of each counterfactual flip.

But we may wonder. What would make it true that the first coin would have landed tails and yet the second would have landed heads? What is it about the coins – or about the rest of the world – that makes the difference? No possible explanation comes to mind. Remember that the coins are perfectly alike. We don't want to say that if all the coins had been flipped, they would have landed the same way – all heads or all tails. The validity of WCN therefore seems to entail that there are inexplicable differences in the counterfactual behaviour of identical coins.

⁴ An assumption made in this argument, which we will not question, is that the conditionals in (4a) and (4b) should receive a univocal analysis. Herburger (2019) denies this, holding that in the context created by 'No student', the relevant conditional means that there is some close p world in which q, whereas in the context created by 'Every student', it means

[Stalnaker 1968] coined the term 'counterfact' for facts that make counterfactual conditionals true. In the case of the coins, the validity of CEM seems to require *primitive counterfacts*. There's a fact of the matter about how the first coin would have landed, about how the second coin would have landed, and so on, and these facts are not grounded in or made true by anything else.

Such primitive counterfacts seem metaphysically objectionable. For one thing, they rule out the "Humean" view that all modal truths are grounded in non-modal truths. They also suggest that current physics is radically incomplete, as it says nothing about how unflipped coins would have landed if they had been flipped. Further, primitive counterfacts raise serious conceptual and epistemological worries. It is hard to see how anyone, even in principle, could come to know how the 100 coins would have landed. No ordinary information seems to shed light on the question. These counterfacts are *inscrutable* from other facts, in the sense of [Chalmers 2012]. There will be little hope for a manageable scrutability base.⁵

One might wonder how a debate over the logic of conditionals – a debate *about words*, about the rules that govern our use of conditional sentences – could have far-reaching implications for topics that don't concern words. Couldn't one accept the validity of CEM and WCN without committing to inscrutable counterfacts?

In formal semantics, truth is commonly defined relative to an "index" or "point of evaluation". A sentence is *valid* iff it is true at every (formally possible) point of evaluation. On this approach, CEM and WCN are valid as long as any point of evaluation makes true exactly one of A>C or A>¬C, given that A is possible. In principle, we could grant that the world doesn't include any primitive counterfacts, but that the "points of evaluation" in the semantics of conditionals do. Many of us, however, would like the points of evaluation in formal semantics to connect with the extra-linguistic world. We would like to know what it takes for an utterance of A > C to be true. To this end, we need to map the formal concept of *truth at a point of evaluation* onto the concept of *truth in an utterance situation*. How could we do this if an utterance situation does not provide the required counterfacts?

[Stalnaker 1980] offered a possible answer: We could say that a sentence is true in an utterance situation iff it is true at *every* point of evaluation compatible with that situation. Other answers can be considered. We could reject the call for truth-conditions: perhaps we don't need to evaluate sentences as true in an utterance situation, as long as we can explain how we reason with them. Alternatively, we might try to weaken the metaphysical implications of CEM by holding that while *there are* primitive counterfacts, it is metaphysically indeterminate *what* they are.⁶

All these views have costs. One obvious problem they all share is that they license speaking *as if* there are primitive counterfacts. All of them accept that if the hundred coins had been flipped then either the second coin would have landed tails or it would have landed heads: either F > H or F > T. Assuming the disquotation schema for truth applies to F > H and F > T, it follows (in classical

that every close p world is a q world. According to this view, (A) and (B) are indeed equivalent, but Premise 2 is false. 5 For a more extended discussion of primitive counterfacts and their philosophical consequences, see [Goodman 2016] and [Hájek 2021].

⁶ This strategy appears in [Cariani 2021]'s discussion of 'will'. Roughly, in order to preserve "Will Excluded Middle" in the face of future openness, Cariani maintains that, given future openness, it is metaphysically indeterminate which future history is the "Thin Red Line". This is akin to saying that though there are primitive future-directed facts, it is indeterminate what they are.

logic) that

(8) Either F > H is true or F > T is true.

But no metaphysically respectable fact could explain the truth of F > H: if it is true, it is a primitive counterfact. Thus:

- (9) a. If F > H is true then there are primitive counterfacts.
 - b. If F > T is true then there are primitive counterfacts.

Combining (8) and (9) we get:

(10) There are primitive counterfacts.

From which, in turn, we can infer the falsity of physicalism etc. If the friend of CEM wants to deny that there are primitive counterfacts, they have to find a fault with this reasoning. They have to reject a seemingly harmless application of classical logic.⁷

We hope we have said enough to motivate double-checking the evidence in support of CEM. Here we focus on the fact that 'would' appears to commute with negation. To assess this evidence, we propose to look at some other operators that display this kind of behaviour. One case in point is 'should'.⁸

4 Should Excluded Middle?

It is easy to hear (11a) and (11b) as saying the same thing.

- (11) a. I don't think that you should leave.
 - b. I believe that you should stay.

Given that 'I don't think that A' is generally interpreted as 'I think that not A' – a point to which we will return – this suggests that 'should' commutes with negation, so that the following schema is valid:⁹

(SCN)
$$\neg$$
Should $A \leftrightarrow$ Should $\neg A$.

SCN is truth-functionally equivalent to the conjunction of "Should Non-Contradiction" (SNC) and "Should Excluded Middle" (SEM):

```
(SNC) \neg(Should A \land Should \neg A).
```

(SEM) Should
$$A \vee Should \neg A$$
.

⁷ They might instead reject the disquotation schema for F > H and F > T. But the argument could be reformulated "in the material mode", without speaking about truth, so this alone wouldn't help.

⁸ Here we are inspired by (but seek to go well beyond) [Todd 2021: 103-6].

⁹ Perhaps the schema, like WCN, should be restricted to "ordinary" propositions A. Seeing as WCN entails SNC (see a few lines down), one might also restrict the hypothesis to non-dilemma cases. Neither restriction substantially affects what we are going to say.

The case for SCN and thereby SEM can be strengthened by looking at negative answers to questions. Intuitively, (12a) and (12b) are equivalent.

- (12) Should I leave?
 - a. No.
 - b. No, you should stay.

As in the case of conditionals, intuitions about these cases are somewhat mixed. Stronger evidence comes from quantified cases: (13a) and (13b) really seem to say the same thing.

- (13) a. No one who is here should leave.
 - b. Everyone who is here should stay.

The argument from (13a) and (13b) to SCN and SEM proceeds much like in section 2 above. The relevant logical form of (13a) and (13b) is plausibly (14a) and (14b), and (14b) is plausibly equivalent to (14c).

(14) a. [No x: Hx] Should Lx

b. [Every x: Hx] Should \neg Lx

c. [Every x: Hx] ¬Should Lx

But if (14c) is equivalent to (13b), then it is hard to deny that \neg Should Lx is equivalent to Should \neg Lx.

Can we conclude that SCN and SEM are valid? This would mean that for every agent S and every conceivable act A, either S should do A or S should refrain from doing A. This seems plainly false. It's easy to think of scenarios in which it is neither the case that you should do A, nor that you should refrain from doing A.

Accepting SCN and SEM as valid would have excessive philosophical costs. As far as we know, nobody has endorsed SEM as valid. Unlike in the case of 'would', the standard view about 'should' is that it doesn't actually commute with negation, despite the evidence of (11)–(13). It *appears* to commute, but this appearance is an illusion.

5 More Excluded Middles

There is a long list of operators that appear to commute with negation even though one would expect that they don't. The list includes, but is not limited to, classical "neg-raisers" such as 'want', 'choose', 'plan', 'think', 'believe', 'expect', 'seem', and 'advise'. (More on neg-raising below.) To illustrate, here is some relevant data about 'think'.

- (15) a. I don't think that Sam is at home.
 - b. I think that Sam is not at home.
- (16) Does Fatima think that Sam is at home?
 - a. No.

- b. No, she thinks he is not at home.
- (17) a. No student thinks they passed.
 - b. Every student thinks they failed.

These pairs appear to be equivalent. Generalising, 'think' appears to commute with negation. This hypothesis, (TCN), would entail "Think Excluded Middle", (TEM):

```
(TCN) \negThink A \leftrightarrow Think \neg A.
```

(TEM) Think $A \vee Think \neg A$.

Think Excluded Middle implies that it is impossible for anyone to lack an opinion about anything: for any subject S and proposition A, either S thinks A or S thinks $\neg A$. How could this be? One might have thought that most two-year olds, for example, have no opinion about (say) whether Homer was a real person. A typical two-year old does not think that Homer was a real person, nor do they think that it's not the case that Homer was a real person.

For another illustration, here is some data about 'want'.

- (18) a. Sam doesn't want to leave.
 - b. Sam wants to stay.
- (19) Does Sam want to leave?
 - a. No.
 - b. No, he wants to stay.
- (20) a. No student wants to leave.
 - b. Every student wants to stay.

Again, these pairs appear to be equivalent, which suggests that 'want' commutes with negation.

```
(WCN) \neg Want A \leftrightarrow Want \neg A.
```

(WEM) Want
$$A \vee Want \neg A$$
.

But if WEM is valid, then it is impossible for anyone to be indifferent towards anything: for any subject S and proposition A, either S wants A to be the case or S wants $\neg A$ to be the case. This is hard to believe.

Next, consider 'usually'. (We only give the strongest, quantificational data from now on, to save space.)

- (21) a. No one is usually here over the break.
 - b. Everyone is usually gone over the break.

These are intuitively equivalent. But the doctrine of *Usually Excluded Middle* is untenable. It would say that any kind of event is either usually the case or usually not the case. But fair coins don't usually land heads, nor do they usually not land heads.

Another operator that seems to fit the present pattern – mentioned as "exactly analogous" to 'would' in [Stalnaker 1980: 95] is the 'in the fiction' operator. Consider:

- (22) a. Nobody who survived in the first season of "Lost" survived in the second season.
 - b. Everybody who survived in the first season of "Lost" died in the second season.

Should we infer that for every piece of fiction and every proposition, either the proposition or its negation is true in the fiction? That the plot of *The Very Hungry Caterpillar* somehow settles exactly how warm it was when the caterpillar ate the pears, and whether one of the caterpillar's descendants will one day be eaten by a spider?¹⁰

We'll mention one last kind of construction that has also been suggested to resemble counterfactual conditionals ([Schlenker 2004]): definite plurals. It is well-known at least since [Fodor 1970] that 'not: the Fs are G' is generally perceived as equivalent to 'the Fs are not G'. We can see this in quantificational environments:

- (23) a. No student's parents are at the ceremony.
- (23) b. Every student's parents are not at the ceremony.

('No student's parents' quantifies over *the parents* of each student.) Here, too, we should not be tempted to accept a principle of "Plural Excluded Middle". If some Fs are G and some are $\neg G$, then 'the Fs are G' is untrue, and so is 'the Fs are not G'. The only obvious way to license Plural Excluded Middle is to declare, as a matter of necessity, that no plurality can have more than one member.

6 Contrary data

We have seen that many linguistic constructions ϕ appear to commute with negation, thereby excluding "middle" cases in which both $\phi(A)$ and $\phi(\neg A)$ are false. Philosophical reflection, however, suggests that such middle cases are possible. In fact, you don't have to be a philosopher to believe that pluralities can have more than one member, that norms can be permissive, or that a two-year old may have no opinions about Homer. These beliefs are common, and they can be expressed in language. This generates data that counts against the relevant excluded middle theses. If 'Sam has no opinions about Homer' is true, how could it be true that Sam either thinks that Homer is real or that he is not real?

We can also explicitly deny instances of excluded middle. Here is such a denial for 'should': "You can eat whatever you want. It's not the case that you should eat pizza, and it's also not the case that you should not eat pizza." This may be a little stilted, but it is a sensible and coherent thing to say. It would be a contradiction if SEM were valid.

¹⁰ Perhaps unsurprisingly, some philosophers have relied on exactly the considerations which many use to support CEM to support "Fiction Excluded Middle" as well; cf. [Blumberg and Holguín 2023].

For all the operators we have reviewed in the previous two sections, the apparent commutativity with negation can be "cancelled" by drawing attention to the possibility that the case at issue might be a middle case in which neither $\phi(A)$ nor $\phi(\neg A)$. If the excluded middle principles were valid, such cancellations would be impossible.

A related piece of linguistic evidence against the excluded middle principles is that the equivalence judgments that support the principles are sensitive to intonational focus. Compare (24a) and (24b).

- (24) a. I don't think that you should eat PIZZA.
 - b. I don't think that you SHOULD eat pizza.

(24a) would normally convey that the speaker thinks that you should not eat pizza, in line with SCN and SEM. (24b) does not. Loosely speaking, the effect of focus under negation is to indicate which part of the negated construction is false. (24a) indicates that you should eat *something other than pizza*. (24b) instead indicates that the connection between you and eating the pizza is not of type *should*. An alternative to *should* is *may*. Thus (24b) draws attention to the "middle" possibility that eating pizza is optional, which cancels the excluded middle inference.

This diagnosis of the focus effect is related to yet another piece of evidence. 'You may A' and 'you should not A' are plausibly contraries: they can't both be true. (25a) and (25b), for example, appear to contradict one another (provided that the two modals are interpreted relative to the same norms).

- (25) a. You may eat pizza.
 - b. You should not eat pizza.

But if 'may' entails 'not should not', then 'may A and may not A' entails 'not should A and not should not A'. By SEM, the latter is a contradiction. But statements like (26) are perfectly coherent.

(26) You may eat pizza and you may refrain from eating pizza.

The equivalence judgements that support SEM therefore clash with other judgments about entailment and validity: with the judgment that (25a) and (25b) are contraries and (26) coherent.

In sum, there are not only philosophical reasons to doubt excluded middle for 'should'. The excluded middle hypothesis also runs counter to linguistic data. We have chosen 'should', but similar data can be generated for the other examples from the previous section.

We contend that it can also be generated for conditionals. We've already seen that plausible metaphysical statements clash with CEM. These statements may not be common among ordinary people, but they look coherent – indeed true – to us. One can also directly cancel CEM by drawing attention to the possibility of "middle" cases. In a wide range of discussions, one finds philosophers expressing the judgement that there is "no fact of the matter" about how a counterfactual chance process would have turned out – for example, about how our 100 coins would have landed. It's neither true that the second coin would have landed heads, the judgment goes, nor that it would have landed tails. (See, among many others, [Adams 1977: p.110], [Jeffrey 1977: p.193], [Lewis 1981: p.331] [Hájek 1997: p.217], [Joyce 1999: pp.172f.], [Van Inwagen 2000: p.14].) Similar judgments are sometimes expressed for counterfactuals with highly unspecific antecedents: if there had been nothing but a large

¹¹ See, for example, [Beaver and Clark 2008] for a more careful treatment of focus that would allow making our point.

metal sphere, what kind of metal would it be? It's neither true that it would be iron, nor that it would be something other than iron. (See, for example, [Lewis 1981: p.329], [Williamson 1988: p.405].) If the doctrine of counterfacts is true, all these judgments are mistaken.¹²

Of course, the validity of CEM does not require the doctrine of counterfacts. According to [Stalnaker 1980], CEM is valid even though there are cases where neither A>C nor A>¬C is true. We've discussed this option in section 3. It is, in our view, the best way to defend CEM. But it still conflicts with linguistic data: in this case, the fact that the inference from CEM to metaphysically sweeping conclusions (as described in section 3) looks valid.

As in the case of 'should', intonational focus helps to draw attention to the middle possibilities, and thereby to undermine or weaken the appearance of excluded middle. Thinking of the 100 coins, one might feel that (26a) and (26b) – parallel to Williams's (4a) and (4b) – are equivalent:

- (26) a. None of the coins would have landed HEADS if it had been flipped.
 - b. All the coins would have landed TAILS if they had been flipped.

This appearance is noticeably weaker if focus is shifted to 'would':

- (27) a. None of the coins WOULD have landed heads if it had been flipped.
 - b. All the coins WOULD have landed tails if they had been flipped.

Finally, as in the case of 'should', 'would' conditionals appear to have a dual in 'might' or 'could' conditionals. (28a) and (28b), for example, appear to be contraries.

- (28) a. If the coin had been flipped, it would have landed heads.
 - b. If the coin had been flipped, it could have landed tails.
 - c. If the coin had been flipped, it could have landed heads.

If they are contraries, however, and CEM is valid, then the conjunction of (28b) and (28c) is a contradiction, which it is not.

To be clear, we are not suggesting that there is decisive linguistic evidence against CEM. Friends of CEM might try to explain away the above data. Perhaps the judgement that there's no fact of the matter about how the coins would have landed is based on a conflation of unassertability with falsity. Perhaps intonational focus on 'would' somehow changes its meaning into 'determinately would'.¹³ Perhaps the apparent duality of 'might' and 'would' conditionals can be explained in some other way. Analogous moves could be made for 'should', 'think', and 'want', if somebody wanted to defend the validity of SEM, TEM, or WEM. All we want to suggest in this section is that there are not just metaphysical, but also linguistic considerations that appear to speak against the excluded middle principles.

¹² The judgement that there's "no fact at the matter" may seem to point towards a three-valued interpretation in which counterfactuals can be true, false, and neither. Most of the cited authors, however, also assert that the relevant counterfactuals are false. The idea is that there is no specific consequent C such that C would be the case if A were the case: the world leaves open what would be the case if A were the case. It's natural to express this in terms of "no fact of the matter". We generally don't find such expressions for other neg-raisers because here it's not the entire world that leaves something open: in "middle cases" for 'should' or 'think', for examples, it's the norms or the subject's beliefs that leave something open.

¹³ This has been suggested to us by xxx.

7 Towards an explanation

We've seen that there is a long list of operators that appear to commute with negation, in much the same way as 'would'. For most of them, accepting the excluded middle principles as valid is not a serious option. There must be some other explanation for the appearance of commutativity. This, we suggest, calls for caution in the case of 'would'. If something looks like cake, one may be tempted to think that it *is* cake, until one notices that it is surrounded by other things that look like cake but are not. (It doesn't matter if the fake cakes are all of the same kind.)

It would be good to know what explains the appearance of commutativity for the other operators. We could then check whether the explanation (or one of the explanations) also works for 'would'. Unfortunately, there is no established answer. The phenomena we have discussed – about 'should', 'think', plurals, etc. – are widely known, and many candidate explanations have been proposed, but none has been widely accepted.

Let's briefly review the main options. For "neg-raisers" like 'think', an old idea is that their apparent commutativity with negation is explained by a divergence between logical form and surface structure: the logical form of 'not thinks A' is 'thinks not A'. The negation is "raised" in the move from logical form to surface structure. (See, e.g., [Fillmore 1963], [Lakoff 1969], [Prince 1976], [Collins and Postal 2014].) To explain the possibility of cancellation, one might assume that the negation can also take genuinely wide scope, so that 'not thinks A' is ambiguous between two logical forms. Conversational context can favour a reading on which the interpreted structure mirrors the surface structure.¹⁴

Could this view be adapted to conditionals? The hypothesis would be that what appear to be negated conditionals are, by default, at the level of interpretation, conditionals with negated consequents. At first glance, this does not look promising. For one thing, it doesn't seem to get a handle on the data involving lexicalized negation and negative answers to questions, as in (2) and (3). It's also unclear how it would help with quantified cases like (4), where syntactic movement seems to be blocked.

- (4) a. No student would have passed if they had goofed off.
 - b. Every student would have failed if they had goofed off.

The bound pronoun 'they' arguably forces the quantifier to take wide scope in (4a).¹⁵ Analogous problems (and others, see [Horn 1978], [Bartsch 1973], [Zeijlstra 2018]), however, also arise for classical neg-raisers like 'think' or 'should'. If a sophisticated form of the syntactic raising hypothesis can be made to work for classical neg-raisers, it might also work for 'would'.

A different direction of explanation goes back to [Bartsch 1973] (see also [Heim 2000]). On this view, use of the relevant operators tends to trigger the associated excluded middle assumption as a presupposition. For definite plurals ('the Fs are G'), this (assumed) presupposition is called a *ho*-

¹⁴ Recent defenders of this view restrict it to a fairly narrow class of cases. [Collins and Postal 2014], for example, suggest that 'think' is a neg-raiser (in this literal sense), but 'should' is not. [Iatridou and Zeijlstra 2013], on the other hand, argue that 'should' is a *positive polarity item* that can't occur in the scope of negation. The only possible interpretation of 'not should p' would therefore be 'should not p'. See [Giannakidou 2011], [Homer 2015] for further discussion. 15 Cf. [Rabern and Todd 2023].

mogeneity presupposition, because the plurality in question is presupposed to be homogenous with respect to the predicate: Either all the Fs are G or all the Fs are not G. [Fintel 1997] and [Schlenker 2004] suggest that conditionals trigger the same kind of presupposition, which would explain the apparent commutativity with negation.

At least in the case of standard neg-raisers, however, we don't seem to be dealing with an ordinary presupposition. If we say of a two-year old (who has never heard of Homer) that he doesn't believe that Homer is a real person, this is intuitively true. Ordinary presupposition failures are commonly judged to be false or defective. (Though this issue is complex, see [von Fintel 2004].) The excluded middle presuppositions for standard neg-raisers also appear to display unusual projection behaviour, as noted in [Gajewski 2005: 68f.], [Romoli 2013], and [Križ 2015]. [Gajewski 2007] suggests that neg-raisers trigger a special kind of "soft" presupposition, as introduced in [Abusch 2005]. [Horn 1989], [Homer 2015] and [Romoli 2013] suggest that the neg-raising effect is best understood as an implicature.¹⁶

We'll briefly outline a particular approach from this family of broadly pragmatic explanations, due to [redacted]. Begin with a positive utterance like (29).

(29) Bob thinks that Sam is smart.

Following the tradition of [Rooth 1996], [Roberts 1996], and [Beaver and Clark 2008], we assume that utterances always address a question. A plausible candidate for the question addressed by (29) is what Bob thinks about Sam's intellectual abilities. Note that this question presupposes that Bob has an opinion about Sam's abilities. Next, assume that negation tends to preserve the addressed question: (30) can still address the question what Bob thinks about Sam's intellectual abilities.

(30) Bob doesn't think that Sam is smart.

Since the accommodated question presupposes that Bob has an opinion, and the literal content of (30) says that Bob doesn't have a favourable opinion, one can infer that he has an unfavourable opinion. The inference is defeasible because (30) could be used to address a different question – for example, whether Bob thinks that Sam is smart.

This approach appears to be applicable just as well to conditionals. We would assume that an utterance of A>C is normally taken to address the question what would be the case if A were the case: whether C would be the case or whether something other than C would be the case. The question presupposes that either C would be the case or \neg C would be the case. If negation tends to preserve the addressed question, an utterance of \neg (A>C) is predicted to convey that \neg C would be the case.

Again, it is not important to our present point whether this line of explanation is correct. Our point is that there are many theories about the mechanisms that explain the apparent commutativity with

¹⁶ If 'would' triggers an excluded-middle presupposition, CEM comes out as *Strawson-valid*: its instances are true whenever their presuppositions are satisfied. On a trivalent ("semantic") approach to presupposition, this implies that no instance of CEM is false. At least for standard neg-raisers, however, the trivalent approach does not look promising, as it appears to make false predictions about projection and about the defeasibility of the neg-raising inference.

¹⁷ Why should this be so? Perhaps it's a brute fact about the information structure of conditionals: the informational focus tends to be on the consequent. Or perhaps its a reflection of the fact that people tend to ignore "middle cases" in which A>C and A>¬C are both false. (See the discussion of hidden variables below.) If we ignore worlds where nothing settles how a coin would land, the only salient possibilities are *it would land heads* and *it would land tails*.

negation of various operators, and that it is worth exploring whether they might apply to the case of conditionals. Given the similarity of the data, we may expect that some of the same mechanisms are at work.

A few steps in this direction have been explored. [Marty et al. 2020], for example, describe empirical tests suggesting that 'would' doesn't behave like a standard presupposition trigger. We welcome this kind of research, but we also want to express a reservation.

What's at issue is the semantic validity of CEM. We conjecture that CEM is not semantically valid: the semantics of conditionals allows for "middle cases" in which A > C and $A > \neg C$ are both false, because all relevant facts about the world leave open whether C would be the case if A were the case. What might be an example? Philosophers who reject CEM often appeal to quantum-mechanical indeterminacy: "No hidden variable" proofs in quantum mechanics seem to show that the present state of the world and the laws don't settle whether, say, a photon will get reflected by a glass pane. If the photon never actually reaches the glass pane, this suggests that there is nothing in the world (not the past, not the laws, and not the future) that could settle whether the photon would have been reflected if it had reached the pane. We seem to have a "middle case". But even here there is room for debate. According to Bohmian mechanics, for example, the outcome *is* settled by a special kind of hidden variable, the pilot wave.

Ordinary people can hardly be assumed to have an opinion on these matters. When presented with a brief vignette of an ordinary situation, they may well believe that the counterfactual outcome is settled by some (undisclosed) facts about the situation. They might even believe that there is always a fact of the matter about what would happen under counterfactual conditions. If so, we predict that their judgements should be in line with CEM. Our conjecture is not that this or that particular instance of CEM is false, or that it is false in such-and-such ordinary circumstances. The conjecture is merely that one can deny CEM without violating the rules of English. 19

8 Conclusion

Subjunctive conditionals appear to commute with negation: we generally treat $\neg(A > C)$ as equivalent to $A > \neg C$. The validity of CEM provides the most obvious explanation of this appearance. In light of the commutativity data, [Williams 2010: 650] argues that the burden of proof is shifted to opponents of CEM. We hope to have shift the burden of proof back to the other side.²⁰

¹⁸ This is also true for unspecific counterfactuals. It isn't obvious, for example, that there is no fact of the matter about whether Bizet and Verdi would have been Italian if they had been compatriots. For all we know, Bizet's parents might have seriously considered moving to Italy before his birth, and one might hold that this would settle the matter.

¹⁹ Compare [Todd 2020, 2021] on 'will'.

²⁰ To be fair, we have not looked at *all* relevant evidence. Some (e.g. [Moss 2013] and [Santorio 2017]) have argued that CEM is supported by judgements about the modal status of conditionals. For example, one might intuit that (31) *might* be true, or that its *probability* is ½, even though opponents of CEM say that it is plainly false.

⁽³¹⁾ If the coin had been flipped it would have landed heads.

This is an interesting argument that deserves a careful response, which we hope to provide in future work.

References

- Dorit Abusch [2005]: "Triggering from Alternative Sets and Projection of Pragmatic Presuppositions"
- Robert Merrihew Adams [1977]: "Middle Knowledge and the Problem of Evil". *American Philosophical Quarterly*, 14(2): 109–117
- Renate Bartsch [1973]: "Negative Transportation" Gibt Es Nicht". *Linguistische Berichte*, 27(7): 1–7
- David I. Beaver and Brady Z. Clark [2008]: *Sense and Sensitivity: How Focus Determines Meaning*. John Wiley & Sons
- Kyle Blumberg and Ben Holguín [2023]: "Fictional Reality", https://philarchive.org/rec/BLUFRG
- Fabrizio Cariani [2021]: The Modal Future. Cambridge: Cambridge University Press
- Fabrizio Cariani and Simon Goldstein [2020]: "Conditional Heresies". *Philosophy and Phenomenological Research*, 101(2): 251–282
- Fabrizio Cariani and Paolo Santorio [2018]: "Will Done Better: Selection Semantics, Future Credence, and Indeterminacy". *Mind*, 127(505): 129–165
- David Chalmers [2012]: Constructing the World. Oxford: Oxford University Press
- Chris Collins and Paul M. Postal [2014]: Classical NEG Raising: An Essay on the Syntax of Negation, vol 67. MIT press
- Charles J. Fillmore [1963]: "The Position of Embedding Transformations in a Grammar". WORD, 19(2): 208–231
- Kai Von Fintel [1997]: "Bare Plurals, Bare Conditionals, and Only". *Journal of Semantics*, 14(1): 1–56
- Janet Dean Fodor [1970]: *The Linguistic Description of Opaque Contents*.. PhD Thesis, Massachusetts Institute of Technology
- Jon Robert Gajewski [2005]: *Neg-Raising: Polarity and Presupposition*. PhD Thesis, Massachusetts Institute of Technology
- [2007]: "Neg-Raising and Polarity". Linguistics and Philosophy, 30(3): 289–328
- Anastasia Giannakidou [2011]: "Positive Polarity Items and Negative Polarity Items: Variation, Licensing, and Compositionality". In Claudia Maienborn, Klaus von Heusinger and Paul Portner (Eds.) Semantics: An International Handbook of Natural Language Meaning, De Gruyter Mouton, 1660–1712

- Jeremy Goodman [2016]: "Consequences of Conditional Excluded Middle"
- Alan Hájek [1997]: ""Mises Redux" Redux: Fifteen Arguments against Finite Frequentism". *Erkenntnis*, 45(2-3): 209–227
- [2021]: "Contra Counterfactism". Synthese, 199(1): 181–210
- Irene Heim [2000]: "Degree Operators and Scope": In Proceedings of SALT X
- James Higginbotham [1986]: "Linguistic Theory and Davidson's Program in Semantics". In Ernest LePore (Ed.) *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Blackwell, 29–48
- Vincent Homer [2015]: "Neg-Raising and Positive Polarity: The View from Modals". Semantics and Pragmatics, 8: 4:1–88
- Laurence R. Horn [1978]: "Remarks on Neg-Raising". In P. Cole (Ed.) *Pragmatics*, chapter Pragmatics. Brill, 129–220
- [1989]: A Natural History of Negation. Chicago: University Of Chicago Press
- Sabine Iatridou and Hedde Zeijlstra [2013]: "Negation, Polarity, and Deontic Modals". *Linguistic Inquiry*, 44(4): 529–568
- Richard C. Jeffrey [1977]: "Mises Redux". In *Basic Problems in Methodology and Linguistics:*Part Three of the Proceedings of the Fifth International Congress of Logic, Methodology and Philosophy of Science, London, Ontario, Canada-1975, Springer, 213–222
- James Joyce [1999]: The Foundations of Causal Decision Theory. Cambridge: Cambridge University Press
- Manuel Križ [2015]: *Aspects of Homogeneity in the Semantics of Natural Language*. PhD Thesis, University of Vienna
- Robin Lakoff [1969]: "A Syntactic Argument for Negative Transportation". In *Proceedings from the Annual Meeting of the Chicago Linguistic Society*, vol 5. Chicago Linguistic Society, 140–147
- David Lewis [1981]: "Causal Decision Theory". Australasian Journal of Philosophy, 59: 5-30
- Matthew Mandelkern [2018]: "Talking About Worlds*". *Philosophical Perspectives*, 32(1): 298–325
- Paul Marty, Jacopo Romoli and Paolo Santorio [2020]: "Counterfactuals and Undefinedness: Homogeneity vs Supervaluations": 21
- Sarah Moss [2013]: "Subjunctive Credences and Semantic Humility". *Philosophy and Phenomenological Research*, 87(2): 251–278

- Ellen F. Prince [1976]: "The Syntax and Semantics of Neg-Raising, with Evidence from French". Language
- Brian Rabern and Patrick Todd [2023]: "A Puzzle about Scope for Restricted Deontic Modals". *Snippets*, 44: 8–10
- Craige Roberts [1996]: "Information Structure in Discourse: Toward a Unified Theory of Formal Pragmatics". *Ohio State University Working Papers in Linguistics*, 49: 91–136
- Jacopo Romoli [2013]: "A Scalar Implicature-Based Approach to Neg-Raising". Linguistics and Philosophy, 36(4): 291–353
- Mats Rooth [1996]: "Focus". In Shalom Lappin (Ed.) *The Handbook of Contemporary Semantic Theory*, Oxford: Blackwell
- Paolo Santorio [2017]: "Conditional Excluded Middle in Informational Semantics". In *Proceedings of the 21st Amsterdam Colloquium*, Institute for Logic, Language and Computation Amsterdam, 385–394
- Philippe Schlenker [2004]: "Conditionals as Definite Descriptions". *Research on Language and Computation*, 2: 417–462
- 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
- Patrick Todd [2020]: "The Problem of Future Contingents: Scoping out a Solution". *Synthese*, 197(11): 5051–5072
- [2021]: The Open Future, Oxford University Press.
- Peter Van Inwagen [2000]: "Free Will Remains a Mystery: The Eighth Philosophical Perspectives Lecture". *Philosophical perspectives*, 14: 1–19
- Kai von Fintel [2004]: "Would You Believe It? The King of France Is Back! (Presuppositions and Truth-Value Intuitions)". In Marga Reimer and Anne Bezuidenhout (Eds.) Descriptions and Beyond, Oxford University Press
- Kai Von Fintel and Sabine Iatridou [2002]: "If and When If-Clauses Can Restrict Quantifiers"
- J. Robert G. Williams [2010]: "Defending Conditional Excluded Middle". Noûs, 44(4): 650-668
- $Timothy\ Williamson\ [1988]:\ "Bivalence\ and\ Subjunctive\ Conditionals".\ \textit{Synthese},\ 75(3):\ 405-421$
- Hedde Zeijlstra [2018]: "Does Neg-Raising Involve Neg-Raising?" Topoi, 37(3): 417-433