

Considering the Exceptions: On the Failure of Cumulative Transitivity for Indicative Conditionals

Ryan Simonelli

0 Introduction

According to existing accounts of indicative conditionals, any argument of the following form is valid:

$$\frac{\varphi \rightarrow \psi \quad (\varphi \wedge \psi) \rightarrow \chi}{\varphi \rightarrow \chi}$$

Here, I present one main counterexample, three auxiliary counterexamples, and a general procedure for generating indefinitely many more counterexamples, to show that there exist invalid arguments of this form. To limit the scope of my discussion, I focus on the ramifications of these cases for the most popular approach to conditionals in natural language: possible worlds semantics. I argue that this data poses serious problems to variably strict accounts of conditionals (Lewis 1973; Stalnaker 1968), as such accounts are structurally unable to accommodate it. Dynamic strict accounts (von Stechow 2001; Gillies, 2007; Willer 2017), however, are a different story. While existing dynamic strict accounts do not accommodate the data, they are in principle able to, and I propose a modified dynamic strict account, drawing from von Stechow (2001), that does. The key modification is this: whereas existing dynamic strict accounts take into account only the effects of conditional *antecedents* in changing the semantic context, the data shows that we must also take into account the effects of conditional *consequents*

in changing the semantic context. After proposing a modified dynamic strict account, I argue that pragmatic alternatives to the semantic solution proposed here of the sort proposed by Moss (2012) and Lewis (2017) in response to data that has previously motivated dynamic strict accounts, falls short for the data presented here.

1 The Case(s)

Suppose you're at a very large party, with hundreds of people, and an open bar where one can get anything at all that one might want to drink. A friend of yours has told you that there is some woman named Maddy at this party. You've never met Maddy, and you don't know anything about her other than that she's at the party. Now, consider the following sentence:

1. If Maddy's drinking a beer, then she's drinking an alcoholic beverage.

(1) seems true. If Maddy's drinking a beer, then she's drinking an alcoholic beverage. Now, consider the following sentence:

2. If Maddy's drinking a beer and she's drinking an alcoholic beverage, then she's not drinking an O'Doul's.

(2) also seems true. If Maddy's drinking a beer and she's drinking an alcoholic beverage, then, since an O'Doul's is a non-alcoholic beverage, she's not drinking an O'Doul's.¹ Finally, consider the following sentence:

3. If Maddy's drinking a beer, then she's not drinking an O'Doul's.

(3) doesn't seem true. It seems that the truth of (3) would rule out the possibility that Maddy's drinking a beer and the beer she's drinking is an O'Doul's. Since

¹If you didn't already know, O'Doul's is a popular non-alcoholic beer.

you don't know anything about Maddy, and she may very well be a teetotaler, such a possibility surely can't be ruled out. Accordingly, (3) is not true. Indeed, we may well say that it's false.

Taking these intuitions at face value, it follows from the truth of (1), the truth of (2), and the falsity of (3), relative to the context you've just considered, that the following argument is invalid:

The Maddy Argument

1. If Maddy's drinking a beer, then she's drinking an alcoholic beverage.
 2. If Maddy's drinking a beer and she's drinking an alcoholic beverage, then she's not drinking an O'Doul's.
- So, 3. If Maddy's drinking a beer, then she's not drinking an O'Doul's.

Abstractly represented, this argument is of the following form:

$$\frac{\varphi \rightarrow \psi \quad (\varphi \wedge \psi) \rightarrow \chi}{\varphi \rightarrow \chi}$$

The principle of inference displayed by this argument schema might be called *Cumulative Transitivity*.² Existing semantic proposals for indicative conditionals tell us that Cumulative Transitivity is a valid principle of inference. So, given the meanings of "If . . . then . . ." sentences, one should be able to reason from the truth of (1) and the truth of (2) to the truth of (3). However, as the Maddy Argument demonstrates, one cannot do that.

This is the basic bit of data with which I'll be working here. For most people, the intuitions of the truth of (1), the truth of (2), and the falsity of (3), when these

²To keep the terminology here consistent with the terminology deployed in discussions of substructural logics (for instance, Makinson (2005) and Brandom (2018)), in which "Cumulative Transitivity" is taken to pick out a *structural* rather than *operational* principle, we might want to call this principle "*Conditionalized Cumulative Transitivity*," but I'll just stick with the shorter name here.

sentences are presented in the order in which I've just presented them, relative to the context I've just specified, are quite strong.³ Still, for certain people—for instance, those for whom non-alcoholic beer comes quickly to mind—this particular example might not work. Perhaps you were one of those people. No matter. Nothing hangs on this particular example. The Maddy Argument is just one example of an argument in which the schema of Cumulative Transitivity fails. Here are three more arguments capable of demonstrating the same general failure:

The Norm Argument:

- 1n. ✓ If Norm gave Maddy a rose, then he gave her a red flower.
- 2n. ✓ If Norm gave Maddy a rose and he gave her a red flower, then he didn't give her a white rose.
- So, 3n. # If Norm gave Maddy a rose, then he didn't give her a white rose.

The Frank Argument:

- 1f. ✓ If Frank is a fish, then he can't walk on land.
- 2f. ✓ If Frank is a fish and he can't walk on land, then he's not a mudskipper.
- So, 3f. # If Frank is a fish, then he's not a mudskipper.

The Bella Argument:

- 1b. ✓ If Bella is a bird, then she flies.
- 2b. ✓ If Bella is a bird and she flies, then she's not a penguin.
- So, 3b. # If Bella is a bird, then she's not a penguin.

I'll leave it as an exercise for the reader to specify, for each of these arguments, a context against which speakers will reliably deem (1) to be true, (2) to be true, and (3) to be false. It's not hard, and the fact that it's not means that there is a

³Or, at least, people who are competent with the relevant vocabulary, who know what an O'Doul's is.

class of arguments exemplifying the schema of Cumulative Transitivity that are, at least intuitively, invalid.

Not only is it not hard to generate these arguments, it also not hard to specify a general procedure for generating them. Take some general kind K (such as *beer*, *rose*, *fish*, or *bird*), instances of which generally have some feature F (such as *being alcoholic*, *being red*, *being unable to walk on land*, or *being able to fly*). Now, find a sub-kind K' (such as *O'Doul's*, *white rose*, *mudskipper*, or *penguin*) that is exceptional with respect to F , such that, while that K s generally have feature F , K' s, while still being K s, have a materially contrary feature F^* (such as *being non-alcoholic*, *being white*, *being able to walk on land*, or *being unable to fly*), which precludes them from having feature F . If you do that, you'll generally have one of these triads, since there will generally be (1) a licit inference from the proposition ascribing K to proposition ascribing F , (2) a licit inference from the conjunction of proposition ascribing K and the proposition ascribing F to the negation of the proposition ascribing K' , but (3) an illicit inference from the proposition ascribing K to the negation of the proposition ascribing K' .

Existing theories of indicative conditionals in natural language semantics, despite their differences, unanimously validate Cumulative Transitivity.⁴ So

⁴This includes accounts of truth-conditional varieties (Lewis 1973; Stalnaker 1968), including even McGee's (1985) semantics that invalidates modus ponens, suppositional varieties (Adams 1966, 1975; Edgington 1995, 2001), informational varieties (Yalcin 2007, 2012, 2016; Bledin 2014), and dynamic varieties (von Fintel 2001; Gillies 2004, 2007, 2009; Willer 2017). Some of these accounts, such as Lewis (1973), von Fintel (2001), and Gilles (2007), are proposed for counterfactuals, but they can be straightforwardly carried over as accounts of indicatives. In what follows, I will talk as if they are simply proposed as accounts of indicatives. Following Rothschild (2020), I do not take Kratzer's (1981, 1986) restrictor view to be a view of conditionals competing with these other views, but, rather, a view of the meaning of "if" that is compatible with multiple views of the meanings of full "If . . . then . . ." sentences, but, in any case, it should be clear that it validates arguments of this form in much the same way the Lewis/Stalnaker account does. Though I do not explicitly consider suppositional views here, just to be clear that such views fall within the target range, note that Adams (1975) describes this inference pattern as "universally probabilistically sound in that the uncertainty of the conclusion can never exceed the sum of the uncertainty of the premises," (22).

these cases potentially spell trouble for a wide range of views. To limit the scope of the discussion, I will restrict my attention in what follows to how different possible worlds analyses of the semantics of conditionals may or may not be prepared to accommodate these cases. My first task will be to show that these failures of Cumulative Transitivity pose a serious problem for the standard truth-conditional semantics for indicative conditionals, owed to Lewis (1973) and Stalnaker (1968).

2 The Standard Truth-Conditional Account

Before getting to the standard truth-conditional semantics, let me briefly go over some familiar territory to be clear how we get there. It is clear that not only counterfactual conditionals but indicatives as well have a certain sort of counterfactual robustness, and so a material conditional analysis, according to which $\varphi \rightarrow \psi$ is true at w just in case φ is false at w or ψ is true at w , will not do. The simplest candidate semantics, then, is the strict conditional analysis, which says that a conditional $\varphi \rightarrow \psi$ is true just in case every possible world in which φ is true is one in which ψ is also true. We might state this condition as follows:

$$\mathbf{S}_{\rightarrow}(\text{strict}) : \llbracket \varphi \rightarrow \psi \rrbracket^{c,w} = 1 \text{ iff } \llbracket \varphi \rrbracket^c \subseteq \llbracket \psi \rrbracket^c$$

The main problem with the strict conditional analysis is that, according to it, almost all conditionals we ordinarily assent to come out false.⁵ For instance, we are forced to say that (1), in the above example is false, since, after all, it's possible that Maddy's drinking an O'Doul's. Given the widespread acceptance

⁵Of course, one might wonder whether that's actually the case. Alan Hájek (2014), for instance, accepts this conclusion for counterfactuals, claiming that the only true counterfactuals are the strictly necessary ones or the ones with impossible antecedents, and one might be persuaded to accept such a view for indicatives as well. I will proceed on the assumption that such a radical stance is not an option for the purposes of natural language semantics.

of indicative conditionals like (1), (1n), (1f), and (1b) that express reasonable but defeasible inferences rather than strict necessities, the strict conditional analysis seems unacceptable. To be precise, we can understand the defeasibility of these conditionals in terms of the fact that the following principle, the *Monotonicity* of conditionals, fails to hold:⁶

$$\frac{\varphi \rightarrow \psi}{(\varphi \wedge \chi) \rightarrow \psi}$$

as illustrated by examples like the following:

1. ✓ If Maddy's drinking a beer, then she's drinking an alcoholic beverage.
- So, 4. # If Maddy's drinking a beer and she's drinking an O'Doul's, then she's drinking an alcoholic beverage.

The correct analysis of indicative conditionals, it seems, must enable us to accommodate the acceptability yet defeasibility of conditionals like (1), enabling us to say that (1) can be true in a context like the one specified above with (4) nevertheless being false. The textbook “variably strict” semantics for conditionals, owed to Lewis (1973) and Stalnaker (1968, 1975), does just that.⁷

The basic idea of the variably strict semantics is that a sentence of the form $\varphi \rightarrow \psi$ is true in a world w and context c just in case, given a relation of “closeness” determined by c , all of the “closest” worlds in which φ is true are worlds in which ψ is true. Now, it's not entirely clear as to how, exactly, this closeness relation should be spelled out for indicative conditionals. One can spell it out in more objective, metaphysical terms—in terms of similarity to the actual world, as it is generally spelled out in discussions of counterfactuals (Lewis

⁶This principle is usually referred to as “Strengthening the Antecedent.”

⁷For textbook presentations, see, for instance, von Stechow and Heim (2011, 63-66), and Dever (2021, 320-327). It's worth noting that Lewis himself only endorsed this semantics for counterfactual conditionals, endorsing a truth-functional analysis of indicatives.

1973, von Stechow 2001)—or in more subjective, epistemic terms—for instance, in terms of an “expectation pattern” for how things in the world should be, given how they normally are (Veltman 1996, Yalcin 2016).⁸ Since it’s really the structural features of the semantics that concern us here, we can abstract from these potential differences in substance, officially supposing just that a context c supplies an ordering relation between worlds, \leq , such that, for each world w , $w' \leq_w w''$ just in case w' is at least as close to w as w'' . We can then define, for each world w , a function, \min_{\leq_w} , that takes a sentence φ , and returns the set of minimally distant φ -worlds, relative to w , as follows:⁹

$$\min_{\leq_w}(\varphi) = \{w' \mid w' \in \llbracket \varphi \rrbracket^c \text{ and, for all } w'', \text{ if } w'' \in \llbracket \varphi \rrbracket^c, \text{ then } w' \leq_w w''\}$$

Having defined such a function, the semantic value of a conditional sentence $\varphi \rightarrow \psi$ is defined as follows:

$$\mathbf{S}_{\rightarrow}(\text{variably strict}) : \llbracket \varphi \rightarrow \psi \rrbracket^{c,w} = 1 \text{ iff } \min_{\leq_w}(\varphi) \subseteq \llbracket \psi \rrbracket^c$$

So a conditional of the form $\varphi \rightarrow \psi$ is true, relative to a context c and world w , just in case the closest φ -worlds, relative to w , are also ψ -worlds.

⁸If one opts for the former sort of rendering of the ordering relation, one will need to suppose some further things about the actual world—for instance, that Maddy’s normal with respect to her beer-drinking habits—in order to say that (1) is actually true. Plausibly, on such an account, speakers judge (1) is true because most people do have normal beer-drinking habits, and so they assume that an ordering relation of this sort is at play in the context specified above. If, on the other hand, one opts for the latter sort of rendering, then (1) is true at the context specified above simply in virtue of the expectation that people will be drinking alcoholic beer. On an epistemic conception of the ordering relation of this sort, it is perhaps best to think of the “actual world,” which figures into the semantics as the minimal element of the order, as a *representation* of the actual world, rather than the actual world itself. Really, then, one will be assigning conditions for the judgment of the truth of sentences—acceptability conditions, rather than truth-conditions per se. The question of which option is to be preferred—whether the semantic machinery here is to be understood in more objective, metaphysical terms or more subjective, epistemic terms—is a foundational one that I’ll bracket for my purposes here.

⁹I make neither the uniqueness nor the limit assumption in this presentation. I’ll rely somewhat on not making the uniqueness assumption in the presentation that follows, but I don’t wish to take a stand on the limit assumption one way or the other.

It is clear that this semantics invalidates Monotonicity. For instance, the closest worlds in which Maddy's drinking a beer will presumably all be worlds in which she's drinking an alcoholic beverage, so (1) is true, but the closest worlds in which she's drinking a beer and the beer she's drinking is an O'Doul's will not be, so (4) is false. What is perhaps less immediately clear is that, for the very same reason, the variably strict semantics invalidates the following principle, which we might call "Simple Transitivity":¹⁰


$$\frac{\varphi \rightarrow \psi \quad \psi \rightarrow \chi}{\varphi \rightarrow \chi}$$

Counterexamples to this principle are perhaps less well-known, but still widely acknowledged in the literature, and there's a simple recipe for turning any failure of Monotonicity involving the sort of defeasibility at issue here into a failure of Simple Transitivity. Consider, for instance, the following example:

1. ✓ If Maddy's drinking a beer, then she's drinking an alcoholic beverage.
 5. ✓ If Maddy's drinking an O'Doul's, then she's drinking a beer.
- So, 6. # If Maddy's drinking an O'Doul's, then she's drinking an alcoholic beverage.

To see how the semantics invalidates this argument in just the way that it invalidates Monotonicity, let b be "Maddy is drinking a beer," a be "Maddy is drinking an alcoholic beverage," o be "Maddy's drinking an O'Doul's," and consider the following diagram:

¹⁰This principle is usually just called "Transitivity" or "Hypothetical Syllogism." For a discussion, see von Fintel and Heim (2011), 64-66.



Here, the closest world to w in which Maddy's drinking a beer is one in which she's drinking an alcoholic beverage, so (1) is true. The closest world in which Maddy's drinking an O'Doul's is a world in which she's drinking a beer, so (5) is true. But (6) is false, since the closest world in which Maddy's drinking an O'Doul's is not one in which she's drinking an alcoholic beverage.

Despite invalidating Simple Transitivity, a variably strict semantics will inevitably validate the principle that we've called *Cumulative Transitivity*:

$$\frac{\varphi \rightarrow \psi \quad (\varphi \wedge \psi) \rightarrow \chi}{\varphi \rightarrow \chi}$$

To see this, note first that, on a truth-conditional conception of meaning, an argument of the form

$$\frac{\varphi \quad \psi}{\chi}$$

is valid just in case, for every world w and context c , if φ is true in w at c , and ψ is true in w at c , then χ is true in w at c . Now consider an arbitrary world w and context c . If $\varphi \rightarrow \psi$ is true in w at c , then the closest φ -worlds, relative to w , are also ψ -worlds. If $(\varphi \wedge \psi) \rightarrow \chi$ is true in w at c , then the closest worlds, relative to w , that are both φ -worlds and ψ -worlds are also χ -worlds. Since the closest φ -worlds are ψ -worlds, and the closest worlds that are both φ -worlds and ψ -worlds are χ -worlds, it follows that the closest φ -worlds are χ -worlds.

So, $\varphi \rightarrow \chi$ is true in w at c . Thus, any argument of the above form is valid. Since the Maddy Argument is an invalid argument of this form, as demonstrated by the fact that we've specified a context relative to which (1) is true, (2) is true, and (3) is false, the standard truth-conditional account fails.

I will consider responses that the truth-conditional semanticist might offer in Section 5. First, however, let me propose a non-truth-conditional way of dealing with the data.

3 The Dynamic Strict Account

In the last few decades, several authors have proposed a dynamic alternative to thinking about meaning solely in terms of truth-conditions, and a new theory of conditionals has emerged from this paradigm.¹¹ The basic idea of a dynamic semantics is this: rather than thinking of the meaning of a sentence solely in terms of the conditions under which it is true, we can think of the meaning of a sentence, at least in part, in terms of its potential, when uttered in a given context, to change (or “update”) that context. In a slogan, the meaning of a sentence is, at least in part, its context change potential. Of course, the stronger, and far catchier, slogan gets rid of the parenthetical “at least in part,” but I do not intend to argue for this stronger slogan here, and I do not think I need to in order to encounter significant resistance from the orthodoxy. Even the weaker slogan marks a radical divergence from the truth-conditional paradigm, bringing aspects of what is normally relegated to the pragmatics—the effect of the utterance of a sentence on a discursive context—into the semantics proper. This, I am going to

¹¹For seminal pieces of dynamic semantics, see Heim (1982), Groenendijk and Stokhof (1991), and Veltmann (1996). Unsurprisingly, there has been some push-back against the dynamic turn from proponents of the truth-conditional paradigm. See, for instance, Dever (2013) and Lewis (2014).

suggest, is precisely what is needed to accommodate our data.

The basic view I'll propose to accommodate our data is based on three intuitive ideas. First, conditionals are evaluated relative to a set of possibilities that are "in view." Second, a conditional is true just in case, for every possibility in the set possibilities in view in which the antecedent holds, the consequent holds. Third, the consideration of some conditionals can function to expand the set of possibilities that one has in view when one evaluates the truth of that conditional. So, conditionals are strict conditionals over a set of accessible possibilities that they themselves have the potential to change. The account of conditionals based on these three ideas is accordingly called the "dynamic strict" account of conditionals, as conditionals are treated as strict conditionals over a dynamically evolving set of possibilities. The basic distinction the version of the dynamic strict account to be proposed here and existing versions of the dynamic strict account (von Fintel 2001, Gillies 2007, Willer 2017) is that existing dynamic strict accounts consider only the effects of conditional *antecedents* in expanding the set of possibilities considered for the evaluation of the conditional, whereas the account proposed here will also consider the effects of conditional *consequents* in expanding the set of possibilities. Before I officially propose the modified account that I will endorse here, let me start by presenting the dynamic strict account as it is endorsed by proponents today and showing how it does not accommodate our data.

There are different formal frameworks in which the dynamic strict account can be presented. For ease of exposition, I'll present it here in the framework proposed by von Fintel (2001), with slight modification, but little hangs on this decision for our purposes here.¹² As with the variably strict account, we once

¹²The framework here differs in detail from the "spheres"-based framework proposed by Gillies

again suppose that a context c supplies an ordering relation between worlds \leq , but we additionally take a context c to include an accessibility function h_c , which takes a world w and selects a corresponding set of accessible worlds, those that we've informally described as those that are "in view" of a participant in that context, which von Fintel calls the "modal horizon." Von Fintel's proposal is that conditionals have a semantics that contains both a dynamic component and a truth-conditional component. First, they have the potential to update the accessibility function by making it select additional worlds, expanding the modal horizon. Specifically, von Fintel proposes that conditionals expand the modal horizon by making the accessibility function select, in addition to the worlds that are already in the modal horizon, the closest worlds in which the antecedent holds. Then, they have strict truth-conditions relative to this updated accessibility function, being true just in case all the antecedent-worlds in the modal horizon are consequent-worlds. Where φ and ψ are neither modals nor conditionals, von Fintel's dynamic strict account can be put as follows:

S_→ (Dynamic Strict):

- a **Context Change Potential:** $h_{c[\varphi \rightarrow \psi]}(w) = h_c(w) \cup \min_{\leq_w}(\varphi)$
- b **Truth-Conditions:** $\llbracket \varphi \rightarrow \psi \rrbracket^{c,w} = 1$ iff $(h_{c[\varphi \rightarrow \psi]}(w) \cap \llbracket \varphi \rrbracket^c) \subseteq \llbracket \psi \rrbracket^c$

On this account, the context change potentials of conditionals enables us to accommodate the same invalidities that the variably strict semantics accommodates, but to do so while retaining a strict analysis of their truth-conditions.


(2007), though they make basically the same predictions. Willer (2013) motivates a more complicated dynamic variant of the spheres framework which is developed in Willer (2017). The basic proposal that I make here is implementable in any of these frameworks, and the question which framework one should opt for hangs on issues outside the scope of this paper. For that reason, I do not take the time to develop this system here so that it is able to accommodate embedded conditionals and conditionals containing modal operators, since I take it that an adequate development will require abandoning the truth-conditional dynamic strict model considered here for a more complex dynamic model, which would distract from the main point of this paper.

Consider again the invalidity of Monotonicity, as demonstrated by the following argument:

1. ✓ If Maddy's drinking a beer, then she's drinking an alcoholic beverage.

So, 4. # If Maddy's drinking a beer and she's drinking an O'Doul's, then she's drinking an alcoholic beverage.

To see how this argument comes out invalid on the dynamic strict account, consider the following diagram:



Here, updating the original context c with $b \rightarrow a$ expands the modal horizon to include the closest b -worlds, and, since all of the b -worlds in this updated modal horizon are a -worlds, $b \rightarrow a$ comes out as true relative to $c[b \rightarrow a]$. However, updating $c[b \rightarrow a]$ with $(b \wedge o) \rightarrow a$ expands the modal horizon to include the closest $(b \wedge o)$ -worlds, and these worlds aren't a -worlds, so $(b \wedge o) \rightarrow a$ comes out false relative to $c[b \rightarrow a][(b \wedge o) \rightarrow a]$.

Proponents of the dynamic strict account have argued that it has significant

virtues over its variably strict alternative.¹³ However, at least as it stands, it does not help in dealing with the data that concerns us here. On the dynamic conception of meaning against which the dynamic strict account is proposed, an argument of the form

$$\frac{\varphi \quad \psi}{\chi}$$

is valid just in case, for every world w and context c , if φ is true in w at $c[\varphi]$, and ψ is true in w at $c[\varphi][\psi]$, then χ is true in w at $c[\varphi][\psi][\chi]$.¹⁴ So, an argument is valid just in case, when the premises are successively considered, the context evolving accordingly, and they are all judged to be true, and the conclusion is judged to be true. It follows from these definitions that any argument of the following form is valid:

$$\frac{\varphi \rightarrow \psi \quad (\varphi \wedge \psi) \rightarrow \chi}{\varphi \rightarrow \chi}$$

Take an arbitrary world w and context c , containing an accessibility function h_c and ordering relation \leq , and suppose that $\varphi \rightarrow \psi$ is true in w and at $c[\varphi \rightarrow \psi]$ and $(\varphi \wedge \psi) \rightarrow \chi$ is true in w at $c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi]$. Since $\varphi \rightarrow \psi$ is true in w and at $c[\varphi \rightarrow \psi]$, all of the φ -worlds in $h_{c[\varphi \rightarrow \psi]}(w)$, which includes all the worlds that $h_c(w)$ includes and, in addition, includes all the closet φ -worlds, are ψ -worlds. Now, updating $c[\varphi \rightarrow \psi]$ with $(\varphi \wedge \psi) \rightarrow \chi$ would function to bring

¹³The main virtue of the dynamic strict account, according to proponents such as von Fintel, Gillies, and Willer, is that it is able to accommodate “Reverse Sobel Sequences.” Proponents of the dynamic strict account, such as Moss (2012) and Lewis (2017), respond that these sequences are to be dealt with pragmatically. We’ll consider how the analogues of these responses fare in response to the data here in Section 5.

¹⁴Note that this is not quite von Fintel’s definition. On von Fintel’s definition, an argument of this form is valid just in case, for every world w and context c , if φ is true at c , and ψ is true at $c[\varphi]$, then χ is true at $c[\varphi][\psi]$, (2001, 142). While this suffices for the data with which von Fintel concerns himself, it will prove crucial in dealing with the data that concerns us here that the updating effects of a sentence are always processed before evaluation of the truth of that sentence.

into view the closest $(\varphi \wedge \psi)$ -worlds, if they aren't already included in $h_{c[\varphi \rightarrow \psi]}(w)$. But, since the closest φ -worlds are already included in $h_{c[\varphi \rightarrow \psi]}(w)$, and all of these worlds are ψ -worlds, the closest $\varphi \wedge \psi$ -worlds *are* already included in $h_{c[\varphi \rightarrow \psi]}(w)$. Accordingly, $h_{c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi]}(w) = h_{c[\varphi \rightarrow \psi]}(w)$. Now, since $(\varphi \wedge \psi) \rightarrow \chi$ is true in w at $c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi]$, all of the $\varphi \wedge \psi$ -worlds in $h_{c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi]}(w)$ are χ -worlds. Finally, we consider $\varphi \rightarrow \chi$ relative to $c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi][\varphi \rightarrow \chi]$. $h_{c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi][\varphi \rightarrow \chi]}(w)$ includes all the worlds that $h_{c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi]}(w)$ includes, and, in addition, includes the closest φ -worlds, but, since these worlds are already included in $h_{c[\varphi \rightarrow \psi]}(w)$, $h_{c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi][\varphi \rightarrow \chi]}(w) = h_{c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi]}(w) = h_{c[\varphi \rightarrow \psi]}(w)$. Now we just check whether all the φ -worlds in the modal horizon are χ -worlds. Since all the φ -worlds are ψ -worlds, and all the worlds that are both φ -worlds and ψ -worlds are χ -worlds, all the φ -worlds are χ -worlds. So, $\varphi \rightarrow \chi$ is true in w at $c[\varphi \rightarrow \psi][(\varphi \wedge \psi) \rightarrow \chi][\varphi \rightarrow \chi]$. Thus, any argument of the above form is valid. Since the Maddy Argument is an invalid argument of this form, the dynamic strict account, at least as it stands, fails.

The way in which the dynamic strict semantics goes wrong in validating the Maddy Argument is quite clear: neither (2) nor (3) actually function to change the modal horizon. Given that (1) is true, there is no mechanism for (2) or (3) to expand the modal horizon. (1) updates the context so that the modal horizon includes all the closest beer-worlds. Since, (1) is true, all of the beer-worlds are alcoholic beverage-worlds, so the modal horizon already includes all the closest beer and alcoholic beverage-worlds. Accordingly, an update with (2) idles, and, since the antecedent of (3) is the same as that of (1), so does an update with (3). Clearly, what needs to happen is that an update (2) or (3) needs to bring into view some O'Doul's worlds, so that (3) will be judged to be false, relative to the updated context against which it ends up being considered. But that's not what

happens on the existing version of the dynamic strict account. Unlike the truth-conditional account, however, which is structurally unable to accommodate the data, the dynamic account can easily be modified to accommodate the data.

4 The Modified Dynamic Strict Account

The data here calls for a modification of the context change potentials of conditionals. Now, I have already said what I take the required modification to be: conditional *consequents* must function to bring possibilities into view, in addition to conditional antecedents. However, given what I've said so far, this conclusion might seem unwarranted. One might think that what is doing the work is not the consequent of (2) and (3), but the conjunctive antecedent of (2).¹⁵ There is, as it turns out, a principled reason to maintain that the conjunctive antecedent of (2) may actually function to bring into view possibilities in which Maddy's drinking an O'Doul's. Suppose the original context includes no possibilities in which Maddy's drinking a non-alcoholic beer. In such a case, the conjunctive antecedent in (2) would violate (the conjunctive analogue of) Hurford's (1974) constraint, the second conjunct being redundant, given the first.¹⁶ If we modify the existing dynamic strict account so that a conditional not only semantically presupposes the possibility of its antecedent but also semantically presupposes the semantic presuppositions of its antecedent, we can plausibly get a context change potential for (2) that invalidates the Maddy Argument.¹⁷ In considering

¹⁵Thanks to Malte Willer for raising this possibility.

¹⁶Hurford's constraint is originally articulated as a constraint on disjunctions, but it is straightforwardly extended to conjunctions.

¹⁷For some additional evidence for this proposal, consider the following sentence:

7. If Maddy's drinking a beer and she's either drinking an O'Doul's or an alcoholic beer, then she's drinking an alcoholic beverage.

After judging (1) to be true, speakers will generally not judge (7) to be true, but that's not what the

(2), the potential violation of Hurford’s constraint forces the hearer to add to the modal horizon the closest possibilities in which Maddy’s drinking a non-alcoholic beer, for it is the elimination of such possibilities that give the second conjunct a function after the first. Since O’Doul’s is among the most popular kinds of non-alcoholic beers, the closest possibilities in which Maddy’s drinking a non-alcoholic beer includes some in which she’s drinking an O’Doul’s, and the presence of these possibilities defeats (3).

I don’t intend to reject this proposal entirely—the conjunctive antecedent of (2) may well bring into view possibilities in which Maddy’s drinking a non-alcoholic beer for the reasons just stated. At the very least, this seems to be a possibility worth exploring.¹⁸ However, this can’t be all that’s going on in the case of the Maddy Argument. Suppose, instead of presenting (1), (2), and (3), relative to the context originally specified, I skipped (2), presenting (1) and then jumping straight to (3). In this case, you would have still judged (1) to be true and (3) to be false. Since (1) and (3) have the same antecedent, it’s got to be the consequent of (3) that is functioning to bring O’Doul’s-worlds into view. Somehow, considering a conditional with the consequent that Maddy’s not drinking an O’Doul’s brings into view worlds in which she is. The minimal modification required to deliver the right results, accordingly, is simply one that makes this the case:

S_→ (Dynamic Strict, Modified):

unmodified dynamic strict account predicts. The closest worlds in which the antecedent is true will be worlds in which the first conjunct is true and the second disjunct of the second conjunct is true, but these world’s aren’t O’Doul’s-worlds. On this proposal, the addition to the modal horizon of O’Doul’s worlds is explained by the fact that a disjunction semantically presupposes the possibility of both of the disjuncts, so O’Doul’s-worlds get added to the modal horizon against which (7) is judged to be false.

¹⁸Moreover, given the story that follows, it’s additionally worth exploring whether conjunctive *consequents* can bring possibilities into view in this way.

- a **Context Change Potential:** $h_{c[\varphi \rightarrow \psi]}(w) = h_c(w) \cup \min_{\leq w}(\varphi) \cup \min_{\leq w}(\neg\psi)$
- b **Truth-Conditions:** $\llbracket \varphi \rightarrow \psi \rrbracket^{c,w} = 1$ iff $(h_{c[\varphi \rightarrow \psi]}(w) \cap \llbracket \varphi \rrbracket^c) \subseteq \llbracket \psi \rrbracket^c$

On this modified proposal, updating a context with a conditional brings into view the closest possibilities in which the antecedent holds, but also the closest possibilities in which the consequent doesn't hold, and a conditional is true, relative to the updated context, if every possibility in view in which the antecedent holds is a possibility in which the consequent holds.

The minimal modification of the dynamic strict account is motivated by the simple idea that to assert a conditional sentence is, in the fully felicitous case, to assert that a *connection* obtains between the antecedent and the consequent: the truth of the antecedent *ensures* the truth consequent.¹⁹ In a possibility-based framework, one can see there to be such a connection between the antecedent and the consequent only if there are some possibilities in view in which the consequent *doesn't* hold, so that these possibilities that can be seen to be *ruled out* by the holding of the antecedent. Only by seeing that the possibilities in which the consequent doesn't hold are excluded from the set of possibilities in which the antecedent does hold can one see that the antecedent and the consequent are connected in the right way for the conditional to be judged to be true. If there are no possibilities in view in which the consequent doesn't hold, then the antecedent cannot be distinguished from any other sentence as one that ensures the truth of the consequent.²⁰ So, the judgment of the truth of a non-trivial conditional requires, if there aren't any possibilities in view in which the consequent doesn't

¹⁹When I speak of "ensurance" here, I mean to be speaking of ensurance relative to a set of salient possibilities, not absolute ensurance relative to the total set of possibilities. Hardly any indicative conditional, I would claim, expresses the latter kind of ensurance.

²⁰I bracket the consideration of mathematical conditionals here, such as "If 48 is divisible by 6, then it is divisible by 3," which is felicitous and has a necessary consequent, since such conditionals pose problems generally for any possibility-based semantics for conditionals rather than posing any specific problem for the semantics proposed here.

hold, that such possibilities be brought into view so that they can be seen to be excluded from the set of possibilities in which the antecedent holds. Now, there may well be motivation for a greater modification of the dynamic strict view, for instance, one according to which a conditional consequent not only brings into view possibilities in which it doesn't hold but also possibilities in which it does hold, but I will not consider such a further modification here.²¹ Since the minimal modification is motivated, and that is enough to resolve our issue, that's the proposal with which I'll settle here.

Let me now explicitly state how this account resolves our issue. Once again, on the dynamic strict account proposed here, an argument of the form


$$\frac{\varphi \quad \psi}{\chi}$$

is valid just in case, for every world w and context c , if φ is true in w at $c[\varphi]$, and ψ is true in w at $c[\varphi][\psi]$, then χ is true in w at $c[\varphi][\psi][\chi]$. With the modification to the dynamic strict account, it now does *not* follow that any argument of the following form is valid:

$$\frac{\varphi \rightarrow \psi \quad (\varphi \wedge \psi) \rightarrow \chi}{\varphi \rightarrow \chi}$$

To see this, consider the Maddy Argument. Let b be "Maddy's drinking a beer," a be "Maddy's drinking an alcoholic beverage," o be "Maddy's drinking an O'Doul's," and s be "Maddy's drinking a soda," and consider the following diagram:

²¹One might, for instance, motivate such a view by drawing inspiration from Starr (2014) and arguing that conditionals presuppose that both their antecedent and consequent are live questions, requiring, for both the antecedent and consequent, possibilities in which it does and does not obtain.



Here, updating the original context c with $b \rightarrow a$ brings into view the closest b -worlds and not- a -worlds, and $b \rightarrow a$ is true, relative to $c[b \rightarrow a]$, since all the b -worlds in $h_{c[b \rightarrow a]}(w)$ are a -worlds.²² Now, updating $c[b \rightarrow a]$ with $(b \wedge a) \rightarrow \neg o$ adds the closest o -world, and $(b \wedge a) \rightarrow \neg o$ is true, relative to $c[b \rightarrow a][(b \wedge a) \rightarrow \neg o]$, since all the $(b \wedge a)$ -worlds in $h_{c[b \rightarrow a][(b \wedge a) \rightarrow \neg o]}(w)$ are $\neg o$ -worlds. Nevertheless, since the updated context now includes an o -world, which is a b -world, it's not the case that all of the b -worlds in $h_{c[b \rightarrow a][(b \wedge a) \rightarrow \neg o][b \rightarrow \neg o]}(w)$ are $\neg o$ -worlds, so, $b \rightarrow \neg o$ is false, relative to $c[b \rightarrow a][(b \wedge a) \rightarrow \neg o][b \rightarrow \neg o]$. Thus, the Maddy Argument is invalid.

5 Pragmatic Alternatives

When I originally presented the Maddy Argument, I first presented a context and then presented three sentences: (1), (2), and (3), in that order. For those with

²²It is worth being clear that, on this proposal, if the original modal horizon doesn't include any possibilities in which Maddy's not drinking an alcoholic beverage, $b \rightarrow a$ functions to bring the closest such possibilities into view. I assume, reasonably I think, that the closest such possibilities are ones in which she's, say, drinking soda, not ones in which she's drinking an O'Doul's.

standard intuitions, (1) is deemed true, relative to the context against which it gets evaluated, (2) is deemed true, relative to the context against which it gets evaluated, and (3) is deemed false, relative to the context against which it gets evaluated. When we considered how the argument is validated by the variably strict semantics, we worked on the assumption that (1), (2), and (3) were all evaluated relative to the same context. We have now presented a semantics according to which the argument is invalid in virtue of the fact that the context does not stay fixed, but, rather, shifts through the process of sequentially evaluating the sentences that constitute it. When the first premise is evaluated, there are no O'Doul's-worlds in view, but the presentation of (2) updates the context by bringing into view O'Doul's worlds, and so, when the conclusion is evaluated relative to the updated context containing O'Doul's-worlds, it is evaluated as false. It seems clear that taking into account how the context changes through the evaluation of these sentences in this sort of way is necessary to understand what is going on in these cases. One may still ask, however, whether this contextual evolution needs to be understood as a matter of the *semantics* of conditional sentences or whether it can be understood as a pragmatic matter, with the standard variably strict semantics for conditionals being maintained. I will consider two attempts to understand this data pragmatically, the first owed to Sarah Moss (2012) and the second owed to Karen Lewis, both proposed in response to the data that has hitherto motivated the dynamic strict account, and argue that they both fall short for the data that concerns us here.²³

Let me start with the sort of account proposed by Moss.²⁴ Suppose, as we

²³The specific cases with which Moss and Lewis are concerned are "Reverse Sobel Sequences" involving counterfactuals, originally observed by Irene Heim and discussed by von Stechow (2001), which have been the main motivation for previous dynamic strict proposals.

²⁴It's worth being clear that, while this account can be straightforwardly carried over for indicatives, Moss herself only endorses it for counterfactuals, endorsing an unorthodox semantics

have been, that the closest worlds in which Maddy's drinking beer are ones in which she's drinking an alcoholic beverage and not an O'Doul's. So, (1) and (2) are true. Now, when you consider (3), after having judged that (1) and (2) are true, it surely seems that (3) is false. Of course, things are not always as they seem, and, if one endorses a pragmatic account of the sort endorsed by Moss, one will say that things are not as they seem in this case: though (3) seems false, it is really true, following from the truth of (1) and (2). The reason it seems wrong to utter (3), relative to the context against which it gets considered, because the presentation of (2) functions bring into view worlds in which Maddy's drinking an O'Doul's, and, once these worlds have been made salient, it's epistemically irresponsible to utter this sentence, given that these worlds can't be ruled out. The key idea is to make a sharp distinction between the *semantic* properties of the truth or falsity of conditionals, which depend only on the ordering relation and are not taken to vary depending on which possibilities have been made salient, and the *pragmatic* properties of the epistemic responsibility or irresponsibility of uttering conditionals, which are taken to vary depending on which possibilities have been made salient. The basic claim of the error theory, meant to explain speakers' judgment that (3) is false as erroneous, is that speakers systematically mistake the pragmatic property of epistemic irresponsibility for the semantic property of falsity.

Accounts along the lines of that proposed by Moss get much of their plausibility from an apparent analogy with other cases in which it seems wrong to utter certain sentences, not because they're false, but because it's epistemically irresponsible to make these utterances once certain possibilities have been raised

for indicatives and other epistemic vocabulary in which semantic values are sets of probability measures (Moss 2015).

to salience. Consider, for instance, the famous case from Dretske (1970). You go to the zoo and see a black and white striped animal in the zebra pen. Intuitively, it seems fine for you to say, "That's a zebra." However, suppose, prior to your saying this, your friend raises the possibility that the animal in the pen might be a cleverly disguised mule. If your friend does this, then, assuming you can't rule this possibility out, it's going to seem wrong to then say "That's a zebra." Since the truth-conditions of "That's a zebra" depend only on whether the animal in the pen is a zebra, and not on which possibilities are salient, the difference here is clearly not a matter of your sentence being true in the first case and false in the second. So, if our semantics is a truth-conditional one, the explanation of why the first utterance seems fine and why the second utterance seems wrong is not going to be a semantic one.²⁵ Rather, it's going to be a pragmatic one—the utterance of "That's a zebra" in the second case seems wrong because, though the sentence is true (assuming there really is a zebra in the pen), it's irresponsible to utter it once the possibility that the animal in the pen is a cleverly disguised mule has been made salient. The proposal is that just the same sort of thing is going on in the cases that concern us here.

While this analogy might seem to help Moss's case, it works only if we are able to maintain that, like the utterance of "That's a zebra" in the zebra case, the possibilities that are contextually salient don't affect the truth value of the sentence uttered and only affect the responsibility of uttering of it. Though

²⁵If our semantics is not truth-conditional, then the explanation may well be a semantic one. For instance, if one goes fully dynamic here, thinking of the meanings of entirely in terms of their context change potential, then one will likely substitute the notion of a sentence's being *true*, relative to a context, with the notion of a sentence's being *supported* by a context, where this means that the sentence is informationally redundant with respect to a context, such that updating that context with the sentence does not change that context. If one then proposes a semantics for epistemic "might"s according to which such expressions function to add possibilities to the context, then one can say that the original context supports "That's a zebra," but the original context, once updated with "That might be a cleverly disguised mule" does not.

one might be able to maintain this for the cases of counterfactual conditionals with which Moss concerns herself, it is much harder to maintain for the cases of indicative conditionals that concern us here. Unlike the truth-conditions of non-modal declaratives like “That’s a zebra,” which do not seem to depend on which possibilities are salient, it seems that the truth-conditions of indicative conditionals really do depend on which possibilities are salient. In this regard, indicative conditionals are more like epistemic modals like “That must be a zebra” than non-modal declaratives like “That’s a zebra.” While the truth of “That’s a zebra” doesn’t vary depending on whether your friend has raised the possibility that it’s a cleverly disguised mule, it seems that the truth of “That must be a zebra” does vary in this way. Your friend’s raising the possibility that the animal in the pen is a cleverly disguised mule seems to result in a context relative to which the sentence “That must be a zebra” is not just irresponsible to utter, but false. Likewise, for the cases that concern us here: the truth of “If Maddy’s drinking a beer, then she’s drinking an alcoholic beverage” seems to vary depending on whether the possibility that Maddy’s drinking an O’Doul’s has been made salient. An account along the lines proposed by Moss, which ignores which possibilities are salient in the evaluation of the truth of an indicative conditional, while likely right for non-modal declaratives, and potentially even right for counterfactuals, just seems wrong for epistemic modals and indicative conditionals.

A different sort of pragmatic account, proposed by Karen Lewis (2017), aims to resolve the above issue with Moss’s account by making contextual salience, still taken to be determined by pragmatic mechanisms, feed directly into the semantics so that the truth-conditions of the conditionals under consideration do vary depending on whether certain possibilities are salient. On Lewis’s account,

as new possibilities become contextually salient through pragmatic mechanisms, the contextually supplied relation of “closeness” that figures into the variably strict semantics changes, with the newly salient possibilities becoming among the closest worlds. So, when the original context gets updated with (3), which mentions the possibility that Maddy’s drinking an O’Doul’s, and the truth of (3) gets considered, relative to this updated context, the ordering relation on worlds that belongs to this context shifts so that worlds in which Maddy’s drinking an O’Doul’s become among the closest worlds, and so (3) comes out false according to the variably strict semantics. An account along these lines maintains that, while the Maddy Argument is valid, since, relative to any context, if (1) and (2) are true, relative to that context, then (3) is also true relative to that context, as the variably strict semantics dictates, in any actual consideration of the truth of (1), (2), and (3), the context shifts through the course of this evaluation, and so (1) and (2) may be true, relative to the context against which they are considered, but (3) may be false, relative to the context against which it gets considered. This basic idea of this approach is to grant, with the dynamic strict theorist, that the semantic properties of sentences change as the discourse context evolves, but to nevertheless maintain, in opposition to the dynamic strict theorist, that the dynamics of discourse evolution belongs squarely in the pragmatics, rather than the semantics.

There are two issues I’d like to raise with an account along these lines. The first and most immediate issue with Lewis’s approach is that, while it is proposed in defense of the variably strict semantics, it actually undercuts the main motivation for the variably strict semantics. Recall, one of the main motivations for the variably strict semantics is that it enables us to maintain that arguments like the following are invalid:

1. ✓ If Maddy's drinking a beer, then she's drinking an alcoholic beverage.
- So, 4. # If Maddy's drinking a beer and she's drinking an O'Doul's, then she's drinking an alcoholic beverage.

If one endorses Lewis's pragmatic account for cases like the Maddy Argument, then it's hard to see why one wouldn't take the same line here, maintaining that this argument is valid but appears invalid in virtue of the fact that the context shifts from the evaluation of (1) to the evaluation of (4). Specifically, the thought would be that, originally, no non-alcoholic beer-worlds are salient, but the presentation of (4) updates the discourse context by making salient possibilities in which Maddy's drinking an O'Doul's, and, when (4) is evaluated, relative to this updated context, it is false. It's hard to see why one who endorses an account along the lines of that proposed by Lewis wouldn't endorse a strict semantics in which indicative conditionals are strict conditionals over the set of contextually salient possibilities and just say this. At the very least, one who endorses an account along the lines of that proposed by Lewis owes an explanation of why this approach should be taken for the Maddy Argument and the other cases that concern us here but not for the argument above. That's the first issue, but this brings us to the second issue, which gets at the root of the problem with pragmatic approaches along the lines of that proposed by Lewis.

I take it that, when it comes to the above argument from (1) to (4), the reason to endorse a variably strict semantics or a dynamic strict semantics over a fully strict semantics is that doing so enables us to maintain that this intuitively invalid argument really is invalid. As a general principle, if our semantic theory says that an intuitively invalid argument is valid, this is, all else being equal, a bad result. Now, if our semantic theory is good in enough other places, we may be willing to

cope with a bad result in this one place, but, in general, the value of the notion of validity defined by a semantic framework directly corresponds to the extent to which it tracks speakers' judgments of intuitive validity, as manifested by their judgments of the truth or falsity of the elements of a series of sentences, expressed relative to an initial context. The problem is that, on Lewis's account, what is actually valid and what is intuitively valid may entirely swing free of one another. Lewis acknowledges that the context may change as the sentences constituting an argument are presented, with new possibilities becoming salient as these sentences are presented, and the truth-values of sentences vary depending on which possibilities have been made salient, and so these truth-values can change as the context evolves. However, the notion of "validity" that is defined in the truth-conditional framework to which Lewis adheres does not take into account contextual evolution, but, rather, is defined in terms of contexts that are supposed to stay fixed. Since, as Lewis acknowledges, contexts rarely do stay fixed, the truth-conditional notion of validity to which Lewis adheres can systematically fail, by Lewis's own lights, to track intuitive validity. This is what we observe in the cases that concern us, but there is no reason to think that failures of this sort are not completely ubiquitous across natural language. If that's so, it's hard to see what good the truth-conditional notion of validity is for the purposes of natural language semantics at all. It seems that our purposes would be much better served by a notion of validity that incorporates the context change potentials of sentences, and so actually tracks judgments of intuitive validity and invalidity which often depend on contextual evolution. That, of course, is just what our dynamic notion of validity does. So, contra Lewis, there is a decisive reason to incorporate the evolution of context into the semantics, rather than treating it as a pragmatic matter.

A third sort of pragmatic response might be worth considering. I have just been working on the assumption that one who endorses the variably strict truth-conditional semantics would have to maintain that the Maddy Argument is valid, and so maintain that either (1) or (2) is false, relative to the original context, or that (3) is true. However, one line of response that has been brought to my attention is to maintain that, in fact, (1) and (2) do not entail (3) because “beer” is ambiguous between (1) and (3).²⁶ In (1), “beer” means specifically *alcoholic* beer, and so it is trivially true that if Maddy’s drinking a beer, she’s drinking an alcoholic beverage. However, in (3), once the possibility of non-alcoholic beer has been introduced through pragmatic factors, the speaker’s interpretation of “beer” shifts so that the extension of “beer” includes both *alcoholic or non-alcoholic* beer. As such, inferring (3) from (1) and (2) would be committing a fallacy of equivocation. Now, if one goes this route, then one would surely want to take it generally, and so one would likewise maintain that “fish” in the Frank argument is ambiguous in (1f) and (3f). But this seems to me to be a wildly implausible thing to maintain. On this account, the default semantic value of terms such as “fish” is not a function that maps each world to the set of *fish* in that world, but a function that maps each world to the set of *normal fish* in that world. But if that were the case, the following dialogue should be felicitous:

Norm: Do you have a pet fish?

Maddy: No, I have a pet mudskipper.

If “fish,” as Norm used it, meant specifically fish that didn’t walk on land, then Maddy would be right to respond negatively to his question if she has a pet mudskipper. But clearly, she should respond positively. She *does* have a pet fish—a mudskipper. Mudskippers are fish. Someone who goes the third route

²⁶This was first suggested to me by Michael Kremer.

is thus forced to denying such trivialities as the claim that, given the default interpretation of “fish” by English speakers, the sentence “Mudskippers are fish” is true, and such a position surely cannot be reasonably maintained.

Now, there are different and potentially more sophisticated versions of this third sort of response, which posit different ambiguities in the sentences that constitute the Maddy Argument.²⁷ Rather than responding to every possible such view, I’ll just quote a widely shared general sentiment expressed by Kripke (1977):

It is very much the lazy man’s approach in philosophy to posit ambiguities when in trouble. If we face a putative counterexample to our favorite philosophical thesis, it is always open to us to protest that some key term is being used in a special sense, different from its use in the thesis. We may be right, but the ease of the move should counsel a policy of caution: Do not posit an ambiguity unless you are really forced to, unless there are really compelling theoretical or intuitive grounds to suppose that an ambiguity really is present, (268).

There does not seem to be any compelling theoretical or intuitive grounds to posit any ambiguity involved in these cases, except, of course, that doing so might save the variably strict semantics. Without any such grounds, we should discount this responses of this third sort.

These three sorts of pragmatic responses seem to exhaust the space of possible responses by the proponent of a truth-conditional variably strict semantics for indicative conditionals, and they all have serious problems. Of course, more could be said in defense of any of these pragmatic responses, and one could dig in one’s heels at a number of places, but I conclude that the best thing

²⁷A different version of this third response, suggested to me by Alex Rausch, still takes the argument to be invalid in virtue of being ambiguous, but takes the ambiguity to not be lexical ambiguity of “beer” but scopal ambiguity of “a,” where the indefinite is understood as an existential quantifier.

to say, in response to the data presented here, is that Cumulative Transitivity really is invalid. So, the truth-conditional variably strict semantics for indicative conditionals semantics must be abandoned for a semantics that invalidates Cumulative Transitivity. I have presented one such semantics. It is worth reiterating, however, that I have limited myself to a specific class of approaches to the semantics of conditionals here, and it may well be that a very different sort of approach to the semantics of conditionals better explains the data than the one adopted here. I'll leave the question of whether that is so and, if it is, what that approach might be, to future discussions.

6 Conclusion

It is widely accepted that natural language indicative conditionals do not validate Monotonicity. However, all existing theories of natural language indicative conditionals validate Cumulative Transitivity. The considerations advanced here suggest that, whenever we have exceptions to Monotonicity involving general rules which have exceptions, we also have exceptions to Cumulative Transitivity. These exceptions to Cumulative Transitivity have, for most of the history of the logic of natural language, gone unnoticed. But, as the study of conditionals in natural language has shown us, we must consider the exceptions.²⁸

²⁸I first stumbled upon these cases working with the Research on Logical Expressivism (ROLE) group, led by Robert Brandom and Ulf Hlobil. I am tremendously grateful to Malte Willer for helping me see the potential significance of these cases for the views discussed here and for his guidance at every stage of the process of writing this paper. I've also benefited from comments from and/or discussions with Ginger Schultheis, Michael Kremer, Jim Conant, Bob Brandom, Alex Rausch, Matt Teichman, and two anonymous referees for this journal.

References

- [1] Adams, Ernest. 1965. "The Logic of Conditionals." *Inquiry* 8: 166-197.
- [2] Adams, Ernest. 1975. *The Logic of Conditionals*. Dordrecht: Reidel.
- [3] Bledin, Justin. 2014. "Logic Informed." *Mind* 123, no. 490: 277-316.
- [4] Brandom, Robert. 2018. "From Logical Expressivism to Expressivist Logics: Sketch of a Program and Some Implementations." *Philosophical Issues* 28, no. 1: 70-88.
- [5] Dever, Josh. 2013. "The Revenge of the Semantics-Pragmatics Distinction." *Philosophical Perspectives* 27: 104-144.
- [6] Dever, Josh. 2021. *Formal Methods in Philosophy of Language*. Draft as of July 21, 2021.
- [7] Dretske, Fred. 1970. "Epistemic Operators." *Journal of Philosophy* 67, no. 24): 1007-1023.
- [8] Edgington, Dorathy. 1995. "On Conditionals." *Mind* 104, no. 414: 235-329.
- [9] Edgington, Dorathy. 2001. "Indicative Conditionals." *The Stanford Encyclopedia of Philosophy*.
- [10] von Fintel, Kai. 2001. "Counterfactuals in Dynamic contexts." In *Ken Hale: A Life in Language*, ed. M. J. Kenstowicz, 123-152. Cambridge, MA: MIT University Press.
- [11] von Fintel, Kai and Irene Heim. 2011. *Intensional Semantics*. MIT Lecture Notes.
- [12] Gillies, Anthony. 2004. "Epistemic Conditionals and Conditional Epistemics." *Noûs* 38: 585-616.
- [13] Gillies, Anthony. 2007. "Counterfactual scorekeeping." *Linguistics and Philosophy* 30: 329-360.
- [14] Gillies, Anthony. 2009. "On the truth conditions of if (but not quite only if)." *Philosophical Review* 118: 325-349.
- [15] Hájek, Alan. 2014. "Most Counterfactuals are False." Unpublished manuscript.

- [16] Hintikka, Jaakko. 1962. *Knowledge and Belief*. Ithaca: Cornell University Press.
- [17] Hurford, James. 1974. "Inclusive or Exclusive Disjunction." *Foundations of Language* 11, no. 3: 409-411.
- [18] Kratzer, Angelika. 1981. "The Notional Category of Modality." *Words, Worlds, and Contexts*, ed. H. J. Eikmeyer and H. Reiser, 38-74. Berlin: De Gruyter.
- [19] Kratzer, Angelika. 1986. "Conditionals." *Chicago Linguistics Society* 22, no. 2:1-15.
- [20] Kripke, Saul. 1977. "Speaker's Reference and Semantic Reference." *Midwest Studies in Philosophy* 2: 255-276.
- [21] Lewis, David. 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- [22] Lewis, Karen. 2014. "Do we need dynamic semantics?" In *Metasemantics: New Essays on the Foundations of Meaning*, ed. A. Burgess and B. Sherman, 231-258. Oxford: Oxford University Press.
- [23] Lewis, Karen. 2018. "Counterfactual Discourse in Context." *Noûs* 52, no. 3: 481-507
- [24] Makinson, David. 2005. *Bridges from Classical to Non-Monotonic Logic*. London: King's College Publications.
- [25] McGee, Vann. 1985. "A Counterexample to Modus Ponens." *The Journal of Philosophy* 82, no. 9: 462-471.
- [26] Moss, Sarah. 2012. "On the Pragmatics of Counterfactuals." *Noûs* 46, no. 3: 561-86.
- [27] Moss, Sarah. 2015. "On the Semantics and Pragmatics of Epistemic Vocabulary." *Semantics and Pragmatics* 8, no. 5: 1-81.
- [28] Rothschild, Daniel. 2020 "A Note on Conditionals and Restrictors." In *Conditionals, Probability, and Paradox: themes from the philosophy of Dorothy Edgington*, ed. L. Waters and J. Hawthorne. Oxford: Oxford University Press.
- [29] Stalnaker, Robert. 1968. "A Theory of Conditionals. In *Studies in Logical Theory*, ed. N. Rescher, 98-112. Oxford: Blackwell.
- [30] Stalnaker, Robert. 1975. "Indicative Conditionals." *Philosophia* 5, no. 3: 269-286.

- [31] Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- [32] Starr, William. 2014. "What 'If'?" *Philosophers' Imprint* 14, no. 10: 1-27.
- [33] Veltman, Frank. 1996. "Defaults in Update Semantics." *Journal of Philosophical Logic* 25: 221-261.
- [34] Willer, Malte. 2013. "Indicative Scorekeeping." In *Amsterdam Colloquium*, 249-256.
- [35] Willer, Malte. 2017. "Lessons from Sobel Sequences." *Semantics and Pragmatics* 10, no. 4: 1-57.
- [36] Yalcin, Seth. 2007. "Epistemic Modals." *Mind*, 116 no. 464: 983-1026.
- [37] Yalcin, Seth. 2012. "A Counterexample to Modus Tollens." *Journal of Philosophical Logic* 41: 1001-1024.
- [38] Yalcin, Seth. 2016. "Modalities of Normality." In *Deontic Modality*, ed. N. Charlow and M. Chrisman, 230-255. Oxford: Oxford University Press.