Negating Conditionals in Bilateral Semantics

Malte Willer

University of Chicago
willer@uchicago.edu

Abstract

A recurring narrative in the literature on conditionals is that the empirical facts about negated *ifs* provide compelling evidence for the principle of Conditional Excluded Middle and sit uncomfortably with a large family of analyses of conditionals as universal quantifiers over possible worlds. I show that both parts of the narrative are in need of a rewrite by articulating a bilateral update semantics for conditionals that distinguishes itself from previous frameworks by giving separate acceptance and rejection conditions for conditionals. The resulting framework shows that CEM is inessential for explaining the empirical facts about negated *ifs* but also how the principle can live happily in a strict analysis of conditionals.

1 The Plot

A recurring narrative in the literature on conditionals is that the empirical facts about negated *ifs* provide strong evidence for the principle of Conditional Excluded Middle (CEM) and sit uncomfortably with an analysis of conditionals as universal quantifiers over possible worlds.

$$\text{(CEM)} \qquad \vDash (\phi > \psi) \vee (\phi > \neg\psi)$$

It is a familiar fact that CEM fails to be valid in a classical universalist setting, since whenever some of the relevant ϕ-worlds make ψ true while others make ψ false, neither $\ulcorner\phi > \psi\urcorner$ nor $\ulcorner\phi > \neg\psi\urcorner$ will turn out to be true. The no less familiar arguments to the conclusion that this is a problem rather than a boon reliably appeal to the empirical fact that natural language *ifs* fail to enter into scope relations that seemingly are available if CEM is rejected—the interaction between conditionals and negation being the key case in point.

The basic observation is that conditionals such as (1a) and (1b) ring equivalent:

(1) a. It's not the case that if John takes the exam, he will pass.
    b. If John takes the exam, he won't pass.

The challenge is to explain why (1a) should entail (1b): given CEM, this is just an instance of disjunctive syllogism. In contrast, the inference is a bit of a puzzler if the truth (or acceptance) of a conditional requires all relevant antecedent-verifying worlds to be consequent-verifying worlds (as in, among many others, Lewis 1973; Kratzer 1986; Veltman 1985), since the existence of a relevant world at which John takes the exam and fails is consistent with the existence of a relevant world at which John takes the exam and passes.

That negation causes trouble in a universalist setting is already noted by Lewis (1973), who ends up rejecting CEM but also bemoans that doing so precludes fully accounting for how conditionals play with negation in natural language. Starting with Stalnaker (1981), CEM enthusiasts have repeatedly pressed the latter point in the literature on conditionals, with a variety of glosses. One recent version of the story puts items that lexicalize negation into the spotlight (see Cariani and Goldstein Forthcoming and Santorio 2017). Consider the following pair:

(2) a. I doubt that if John takes the exam, he will pass.
b. I believe that if John takes the exam, he won't pass.

(2a) and (2b) have the air of equivalence. That makes good sense if CEM is true, for then rejecting the claim that John will pass if he takes the exam immediately amounts to accepting the claim that he will not pass if he takes the exam. But it is not at all obvious how the equivalence can be explained if CEM is rejected. Other natural language data suggesting that negated conditionals entail their corresponding conditional negations involve the interaction between quantifiers and conditionals (von Fintel and Iatridou 2002; Higginbotham 2003) and the focus sensitive expression *only* (von Fintel 1997).

This paper demonstrates, first, that CEM is inessential for explaining the empirical facts about negated *ifs* but also, second, that CEM can live happily in a strict analysis of conditionals. I do so by developing a bilateral update semantics for conditionals that distinguishes itself from previous frameworks by giving separate acceptance and rejection conditions for conditionals.

2 Negating Conditionals

I couch my analysis in a bilateral semantic setting that distinguishes between two distinct foundational semantic concepts, for instance: truth-making and false-making, support and rejection, or affirmation and denial. Here I will choose an update-based route and recursively define a positive, acceptance inducing update function $[\cdot]^+$ and in addition a negative, rejection inducing update function $[\cdot]^-$ on a set of possible worlds s.[1] The target language $\mathcal{L}$ is the smallest set that contains a set of propositional atoms $\mathcal{A} = \{p, q, r, \ldots\}$, and is closed under negation ($\neg$), conjunction ($\wedge$), disjunction ($\vee$), and the conditional connective ($>$). I begin by stating the bilateral approach and its immediate payoff (Section 2.1) and conclude by addressing a foundational issue about the analysis of negation proposed here (Section 2.2).

2.1 Framework

Start with the obvious entries for atomic sentences and negation.

$$(\mathcal{A}) \quad \begin{array}{l} s[p]^+ = \{w \in s \colon w(p) = 1\} \\ s[p]^- = \{w \in s \colon w(p) = 0\} \end{array} \qquad (\neg) \quad \begin{array}{l} s[\neg\phi]^+ = s[\phi]^- \\ s[\neg\phi]^- = s[\phi]^+ \end{array}$$

A positive update with p eliminates from the input all possible worlds at which p is false, while a negative update eliminates all possible worlds at which p is true. A positive update with $\ulcorner\neg\phi\urcorner$ is a negative update with ϕ; a negative update with $\ulcorner\neg\phi\urcorner$ is a positive update with ϕ.

Conjunction is commonly modeled as sequential updating in dynamic semantics, but here we do not need to worry about internal dynamics, and so a simple intersective approach à la Veltman 1996 will do. A negative update with a conjunction amounts to taking the union of (i) the result of updating with the negation of the first conjunct and (ii) the result of updating with the negation of the second conjunct. Disjunction receives its standard definition in terms of negation and conjunction.

$$(\wedge) \quad \begin{array}{l} s[\phi \wedge \psi]^+ = s[\phi]^+ \cap s[\psi]^+ \\ s[\phi \wedge \psi]^- = s[\phi]^- \cup s[\psi]^- \end{array} \qquad (\vee) \quad \begin{array}{l} s[\phi \vee \psi]^+ = s[\phi]^+ \cup s[\psi]^+ \\ s[\phi \vee \psi]^- = s[\phi]^- \cap s[\psi]^- \end{array}$$

[1] Thus the upcoming proposal is in the spirit of Veltman 1996 but defines two update functions instead of just one. For some alternative approaches to the bilateral program in semantics, see, among others, Aloni ms.; Fine 2017; Francez 2014; Incurvati and Schlöder Forthcoming; Rumfitt 2000.

More interesting is the analysis of conditionals. Ramsey (1931) famously suggested that a conditional is to be accepted in case its consequent is (hypothetically) accepted under the assumption of the antecedent. Virtually every semantic analysis of conditionals owes inspiration to Ramsey's suggestion, and so does the story I am about to tell here. The key idea of this paper, however, starts with the intuition that there is something important missing from Ramsey's proposal. In saying what it takes to accept a conditional, Ramsey of course also specifies what it takes to *fail* to accept a conditional: do so if you fail to (hypothetically) accept the consequent under the assumption of the antecedent. But he does not—or at least not obviously—state what it takes to *reject* a conditional, for the plain reason that failing to accept a conditional is not trivially the same as rejecting it. It is exactly this omission that I suggest leaves room for an implementation of the Ramsey test for conditionals that lives happily with the empirical facts about negated conditionals. Specifically, I suggest that Ramsey's dictum about accepting conditionals is compatible with the following proposal about what it takes to *reject* a conditional: do so if you (hypothetically) reject the consequent under the assumption of the antecedent.

To make the proposal outlined in the previous paragraph more precise, say that a state s *accepts* ϕ, $s \Vdash^{+} \phi$, just in case $s[\phi]^{+} = s$; and that s *rejects* ϕ, $s \Vdash^{-} \phi$, just in case $s[\phi]^{-} = s$. The basic proposal would then be that a conditional of the form $\ulcorner \phi > \psi \urcorner$ is accepted by s just in case $s[\phi]^{+} \Vdash^{+} \psi$ and rejected just in case $s[\phi]^{+} \Vdash^{-} \psi$. One minor wrinkle: following standard protocol (see, e.g., von Fintel 2001; Gillies 2007; Willer 2017) we shall assume that conditionals presuppose that their antecedent be compatible with the relevant modal domain and since—for current purposes anyway—conditionals impose tests on the input context s, this is just to presuppose that the conditional antecedent is a possibility in the input state. Putting all of this together, we say:

$$(>) \quad \begin{array}{l} s[\phi > \psi]^{+} = \{w \in s \colon s[\phi] \Vdash^{+} \psi\} \text{ iff } s[\phi]^{+} \neq \varnothing \\ s[\phi > \psi]^{-} = \{w \in s \colon s[\phi] \Vdash^{-} \psi\} \text{ iff } s[\phi]^{+} \neq \varnothing \end{array}$$

Positive and negative updates with conditionals both require that their antecedent be compatible with the input state. If defined, a positive update with a conditional tests whether its consequent is accepted once the input state is strengthened with the antecedent. If defined, a negative update with a conditional tests whether its consequent is rejected once the input state is strengthened with the antecedent. A passed test returns the original state; a failed test results in the absurd state ($\varnothing$).

Bilateral setups such as the one we are exploring here allow for a variety of notions of logical consequence, and we will exploit this flexibility momentarily (and use subscripts to keep track of the options). For now, one way to go is to think of valid inferences as those whose premises induce acceptance of the conclusion. Again following standard protocol, we take presupposition failures as constraining the states of information that we have to consider in evaluating an argument for validity: only states for which updating with the premises and the conclusion is defined matter when we check for validity (see e.g. Beaver 2001).

A sequence $\phi_1, \ldots, \phi_n$ *induces acceptance* of ψ, $\phi_1, \ldots, \phi_n \vDash_1 \psi$, just in case for all s such that $s[\phi_1]^{+} \ldots [\phi_n]^{+}[\psi]^{+}$ is defined, $s[\phi_1]^{+} \ldots [\phi_n]^{+} \Vdash^{+} \psi$.

Thinking of entailment as acceptance inducing is to recreate update-to-test consequence (with the possibility of presupposition failures) in a bilateral setting. A special case: say that a state s *admits* ϕ just in case updating s with ϕ (positively or negatively) is defined. Then ϕ is a *validity* just in case every state that admits ϕ also accepts ϕ.

The resulting framework then predicts that negated conditionals entail conditional negations (and vice versa):

Fact 1 $\neg(\phi > \psi) \vDash_1 \phi > \neg\psi$

This is easy to see since every positive update with $\ulcorner\neg(\phi > \psi)\urcorner$ is a negative update with $\ulcorner\phi > \psi\urcorner$, which is testing whether ψ is rejected—and thus whether $\ulcorner\neg\psi\urcorner$ is accepted—under the assumption of ϕ. But this is just what $\ulcorner\phi > \neg\psi\urcorner$ is asking, which establishes the fact.

At the same time, CEM turns out to be invalid if validity amounts to guaranteed acceptance by every admitting state:

Fact 2 $\nvDash_1 (\phi > \psi) \vee (\phi > \neg\psi)$

A very simple counterexample is a state that accepts p but is agnostic about q, say $s = \{w_1, w_2\}$ such that $w_1(p) = w_2(p) = 1$ and $w_1(q) = 1$ while $w_2(q) = 0$. Then $s[p]^+ = s$ but clearly $s \nVdash^+ q$ and $s \nVdash^- q$ and hence $s[p > q]^+ = s[p > \neg q]^+ = \varnothing$. So $s[(p > q) \vee (p > \neg q)]^+ = \varnothing$ and since s is not empty we have $s[(p > q) \vee (p > \neg q)]^+ \neq s$ and so $s \nVdash^+ (p > q) \vee (p > \neg q)$, which establishes the fact.

This concludes the first key lesson of the current exercise: that there need not be a trade-off between signing up for a universalist semantics that rejects Conditional Excluded Middle and accounting for how conditionals play with negation in natural language. What allows us to see that—if we adopt a Ramseyan perspective anyway—is the realization that an account of what it takes to *accept* a conditional lives happily with the existence of a separate characterization of what it takes to *reject* a conditional. Once we couch our analysis of conditionals in a *bilateral* setting that is grounded in two primitive basic semantic concepts—truth and falsity, acceptance and rejection, or positive and negative updating—it becomes straightforward to see that the problem with negating conditionals is really about getting the meaning of *negation* straight rather than coming up with a theory of conditionals that validates CEM.

2.2 Frege-Geach Style Worries

Bilateral stories like the one told here inevitably raise the spectre of the Frege-Geach problem. Textbook metaethical expressivism starts with the natural intuition that a sentence like (3a) "Stealing is wrong" is used to express the attitude of disapproval of stealing.[2] Natural as this may sound, it raises the question of what attitude the negation of (3a), i.e. (3b), expresses. Consider:

(3) a. Stealing is wrong.
b. Stealing is not wrong.
c. Not stealing is wrong.

The initial observation here is that (3b) cannot express disapproval of not stealing, since this is what (3c) expresses; nor can it simply express the absence of the attitude of disapproval of stealing, since someone who is thoroughly agnostic or undecided about the moral status of stealing fails to disapprove of stealing but does not endorse (3b). The obvious conclusion to draw here is that an expressivist analysis of the language of morals must appeal to a second basic attitude besides disapproval—say, the attitude of *tolerance*—and that (3b) expresses that very attitude toward stealing.

The challenge then is to explain why (3a) and (3b) seem to express *incompatible* attitudes. This is especially pressing for expressivists, who aim at explaining the inconsistency of sentences

[2]See Schroeder 2010 and references therein for a discussion of metaethical expressivism and its history.

in terms of an incompatibility between the states of mind expressed by these sentences: so (3a) and (3b) must be inconsistent *because* they express incompatible states of mind. But the issue is perfectly general, and the problem, as Schroeder (2008) explains, is that all that the story sketched so far delivers is that (3a) and (3b) express distinct attitudes toward the same kind of action that—as far as the analysis is concerned—might very well be logically unrelated. If all we have are two basic attitudes—tolerance and disapproval—the incompatibility between the states of mind expressed by (3a) and (3b) remains a matter of pure stipulation.

Some have insisted that there is nothing wrong with grounding one's semantics in a primitive notion of incompatibility between basic attitudes (see e.g. Gibbard 2013). My goal here is not to referee this issue or to dive into a detailed discussion of the language of morals but simply to highlight that the current framework has no need for such maneuvers. It would, on first sight anyway, be dissatisfying if we had to *stipulate* that no single consistent state of information can both accept and reject a single sentence at the same time. Fortunately, however, it is easy to verify that the attitudes of acceptance and rejection are incompatible in the following sense:

Fact 3 For all s and ϕ: if $s \Vdash^{+} \phi$ and $s \Vdash^{-} \phi$, then $s = \varnothing$.[3]

No non-absurd state can both accept and reject a sentence. For instance, suppose that s is non-absurd and accepts the conditional "If Mary is in Chicago, then Jack is in Rome" ($c > r$). Then $s[c]^{+}$ is a non-absurd state and, moreover, $s[c]^{+} \Vdash^{+} r$. Since no non-absurd state can both accept and reject that Jack is in Rome (that is, exclusively consist of r-worlds and of $\neg r$-worlds), it follows that $s[c]^{+} \nVdash^{-} r$ and so that $s \nVdash^{-} c > r$, that is, s fails to reject "If Mary is in Chicago, then Jack is in Rome." For parallel reasons, if s is non-absurd and rejects "If Mary is in Chicago, then Jack is in Rome," then the state must fail to accept the conditional.

The previous result depends, of course, on the details of the semantics of negation: we *could* have provided negative entries for some or all of our connectives (or atomic sentences) that make the attitudes of acceptance and rejection compatible with each other. Nonetheless, the fact remains that there is nothing *per se* dubious about a framework that appeals to two primitive semantic notions in semantic theorizing, and that the incompatibility between the attitudes of acceptance and rejection is not a mere matter of stipulation but something that follows from our semantics. Let us explore the framework a bit further.

3 Toward Conditional Excluded Middle

While the story told so far has taught us something important about conditionals—namely, that the challenge of validating Conditional Excluded Middle is separate from the task of getting the facts about negated *ifs* in natural language straight—there is good reason to think that more needs to be said about the interplay between *if*, *or*, and *not*. To see this, consider again the case of a perfectly fair coin flip:

(4) a. If Maria flipped the coin, it landed tails.
b. If Maria flipped the coin, it landed heads.

The simple intuition here is that it makes good sense to assign .5 probability to (4a) and .5 probability to (4b); since (4a) and (4b) are mutually exclusive, their disjunction has a probability of 1 and thus the distinct air of a validity after all (cf. Santorio 2017). So while the

[3] The key underlying observation here is that if defined, the intersection of $s[\phi]^{+}$ and $s[\phi]^{-}$ is guaranteed to be empty (the proof is straightforward via induction). But if $s \Vdash^{+} \phi$ and $s \Vdash^{-} \phi$, then $s[\phi]^{+} = s[\phi]^{-} = s[\phi]^{+} \cap s[\phi]^{-} = s$, and so it follows that s must be the absurd state.

empirical data about negated conditionals do nothing to demonstrate that CEM is valid, other considerations might give it some (if perhaps not irresistible) appeal. In this section I will show how to accommodate this point, and I will do so by exploiting the fact (already indicated earlier) that a bilateral semantics provides us with some interesting options for how to think about logical consequence.

Start again with the difference between acceptance failures and rejection. Earlier we saw why a conditional like "$(p > q) \vee (p > \neg q)$" need not be accepted by an information carrier: if one is certain about p but agnostic about q, one would not accept "$p > q$" and one would not accept "$p > \neg q$," and so the whole disjunction would not be accepted either. But again failing to accept ϕ is not the same as rejecting ϕ, and in particular we can observe that while a state s need not accept "$(p > q) \vee (p > \neg q)$," it cannot reject it either. For doing so requires that $s[(p > q) \vee (p > \neg q)]^- = s$ and so $s[p > q]^- \cap s[p > \neg q]^- = s$, which in turn requires that s must reject both "$p > q$" and "$p > \neg q$." And that, in turn, requires that $s[p]^+$ is consistent but rejects both q and its negation, which simply cannot happen. The case, indeed, generalizes:

Fact 4 For all s, ϕ and ψ: $s[(\phi > \psi) \vee (\phi > \neg\psi)]^- = \varnothing$.

This is an immediate consequence of the fact that no non-absurd state can accept both ψ and its negation (recall Fact 3).

The previous fact is suggestive, since it highlights an alternative way of thinking about logical consequence and thus of validities. Earlier we suggested that an argument is valid just in case updating with the premises rationally commits one to accepting the conclusion; but instead we may think of an argument as valid just in case one could not rationally update with the premises and also reject the conclusion. More precisely:

> A sequence $\phi_1, \ldots, \phi_n$ *excludes rejection* of ψ, $\phi_1, \ldots, \phi_n \vDash_2 \psi$, just in case for all s such that $s[\phi_1]^+ \ldots [\phi_n]^+[\psi]^+$ is defined, $s[\phi_1]^+ \ldots [\phi_n]^+[\psi]^- \not\Vdash \bot$.

Here "$\bot$" represents an arbitrary contradiction. This proposal then treats ϕ as a validity just in case $s[\phi]^- = \varnothing$ for all states that admit ϕ.

In a classical setup, it does not matter whether we say that the premises of an argument induce acceptance of a conclusion or that they exclude its rejection, for whenever $\phi_1, \ldots, \phi_n$ entails ψ, then $\phi_1, \ldots, \phi_n, \neg\psi$ entails $\bot$, and vice versa. But it does matter in the current setting. Specifically, it follows immediately from Fact 4 that Conditional Excluded Middle is valid if entailment requires that updating with the premises exclude rejection of the conclusion.

Fact 5 $\vDash_2 (\phi > \psi) \vee (\phi > \neg\psi)$

If no consistent state can accept ψ and its negation, it cannot reject ψ and its negation either. So in particular, if consistent, $s[\phi]^+$ cannot reject ψ and its negation, and so (if defined in the first place) $s[\phi > \psi]^- = \varnothing$ or $s[\phi > \neg\psi]^- = \varnothing$.

Validity understood as exclusion of rejection ("validity$_2$" for short) is thus an interesting alternative to the more familiar concept of validity as guaranteed acceptance ("validity$_1$" for short). It also preserves the earlier made prediction about the interplay between negated conditionals and conditional negations, since whenever an argument is valid$_1$, it is also valid$_2$:

Fact 6 If $\phi_1, \ldots, \phi_n \vDash_1 \psi$, then $\phi_1, \ldots, \phi_n \vDash_2 \psi$.

No state that accepts ψ can be updated with the negation of ψ without resulting in the absurd state. So whenever $\phi_1, \ldots, \phi_n$ guarantees acceptance of ψ, it also excludes rejection of ψ. Thus validity$_2$ preserves what we want about negated conditionals but also delivers CEM as valid.

The second important upshot of this discussion, then, is that Conditional Excluded Middle lives happily in a framework that treats conditionals as universal quantifiers over a contextually determined set of possible worlds. The key idea here, again, is that rejecting a conditional is not the same as just failing to accept it. This, recall, allowed us to get the facts about negated conditionals straight without endorsing CEM. But it also turns out that while not every instance of CEM needs to be accepted, all of them are resistant to rejection, and this inspires a notion of validity as excluded rejection that delivers CEM as a tautology after all. Let me highlight another notable consequence before I conclude.

CEM is not the only principle that has some appeal but fails to hold in a large variety of conditional analyses. Another is Simplification of Disjunctive Antecedents (SDA):

$$\text{(SDA)} \qquad (\phi \vee \psi) > \chi \vDash (\phi > \chi) \wedge (\psi > \chi)$$

While the case for CEM tends to be indirect, SDA has the air of an intuitive validity:

(5) If you pay cash or with a debit card, you will receive a five percent discount.
$\rightsquigarrow$ If you pay cash, you will receive a five percent discount.
$\rightsquigarrow$ If you pay with a debit card, you will receive a five percent discount.

Intuitively, (5) suggests that you get five percent if you pay cash, *and* if you pay with a debit card. And indeed SDA is valid$_2$ in the framework developed here:[4]

Fact 7 $(\phi \vee \psi) > \chi \vDash_2 (\phi > \chi) \wedge (\psi > \chi)$

The underlying fact is that if χ is accepted in some state s, then χ it is not rejected by any consistent strengthened state $s' \subseteq s$. So if $s[\phi \vee \psi]^+ \Vdash^+ \chi$, then both $s[\phi]^+ \nVdash^- \chi$ and $s[\psi]^+ \nVdash^- \chi$. It follows that neither $\ulcorner \phi > \chi \urcorner$ nor $\ulcorner \psi > \chi \urcorner$ can be rejected by s.

Cariani and Goldstein (Forthcoming) show that, taken together, CEM and SDA have the potential of bringing a number of unwelcome consequences in their wake, given minimal additional assumptions. The perhaps most striking one is that the following principle turns out to be a validity, which Cariani and Goldstein label the Interconnectivity of All Things (IAT):

$$\text{(IAT)} \qquad \vDash ((\phi > \chi) \wedge (\psi > \chi)) \vee ((\phi > \neg\chi) \wedge (\psi > \neg\chi))$$

What makes IAT so unattractive is that it clashes with the uncontroversial fact that one may endorse the conjunction of (6a) and (6b):

(6) a. If Maria comes to the party, it will be fun.
b. If Bill comes to the party, it won't be fun.

So IAT is no good. But (as Cariani and Goldstein lay out in more detail) if CEM is a validity, then so in particular is $\ulcorner((\phi \vee \psi) > \chi) \vee ((\phi \vee \psi) > \neg\chi)\urcorner$. And if such conditionals simplify—as they should if SDA is really valid—it seems to follow right away that IAT is a validity as well.

And yet it is easy to verify that IAT is not a validity in the bilateral system that we have been exploring so far. Consider the instance "$((p > r) \wedge (q > r)) \vee ((p > \neg r) \wedge (q > \neg r))$" and let $s = \{w_1, w_2\}$ with the following distribution of truth-values:

[4] SDA is also valid$_1$ but for current purposes we will be specifically interested in a notion of validity that delivers both CEM and SDA.

	p	q	r
w_1	T	F	F
w_2	F	T	T

It is easy to see that $s \Vdash^+ p > \neg r$ and so $s \Vdash^- p > r$ and, moreover, $s \Vdash^+ q > r$ and so $s \Vdash^- q > \neg r$. Accordingly, $s[(p > r) \wedge (q > r)]^- = s$ and $s[(p > \neg r) \wedge (q > \neg r)]^- = s$ and so $s \Vdash^- ((p > r) \wedge (q > r)) \vee ((p > \neg r) \wedge (q > \neg r))$.

This is a surprising result: if CEM is valid and conditionals of all stripes simplify, how could IAT not be a validity as well? The answer is that validity$_2$ fails to be transitive.

Fact 8 For some ϕ, ψ, and χ: $\phi \vDash_2 \psi$ and $\psi \vDash_2 \chi$ but $\phi \nvDash_2 \chi$.

Specifically, we have, with "$\top$" being any tautology: (i) $\top \vDash_2 ((p \vee q) > r) \vee ((p \vee q) > \neg r)$ and (ii) $((p \vee q) > r) \vee ((p \vee q) > \neg r) \vDash_2 ((p > r) \wedge (q > r)) \vee ((p > \neg r) \wedge (q > \neg r))$, but also (iii) $\top \nvDash_2 ((p > r) \wedge (q > r)) \vee ((p > \neg r) \wedge (q > \neg r))$. The reason: (i) says that no state may consistently reject a certain CEM-instance and (ii) says that once some CEM-instance is accepted, the corresponding IAT instance cannot be rejected. But just because s cannot consistently reject ϕ does not mean that s must accept ϕ. Accordingly, just because some CEM-instance cannot be rejected does not mean that it must be accepted, and so certain IAT-instances may be rejected by non-absurd states of information.

Thinking of validity as exclusion of rejection, then, shows that CONDITIONAL EXCLUDED MIDDLE can live happily in Ramseyan setting, and that it can co-exist with SIMPLIFICATION OF DISJUNCTIVE ANTECEDENTS without negative side effects. More should—and can—be said about the intuitive foundations of this notion of validity; for now I submit that it is worthy of serious attention.

4 Conclusion and Outlook

The framework proposed here—like others before—has taken the Ramsey test as a source inspiration for its story about conditionals, but it adds to the classical Ramseyan question of what it takes to accept a conditional the one of what it takes to reject a conditional. The resulting bilateral semantic analysis of conditionals allows us to disentangle the question of how conditionals play with negation in natural language from the question of whether CONDITIONAL EXCLUDED MIDDLE is valid. It also demonstrates that the principle is consistent with the spirit of an analysis of conditionals as universal quantifiers over a contextually provided set of possible worlds. Other conditional heresies such as SIMPLIFICATION OF DISJUNCTIVE ANTECEDENTS may be adopted without running into triviality results.

Some important questions must be left to another day. I have not said how conditionals can induce non-trivial changes on the common ground if they articulate tests (in brief, by modeling the common ground as a set of sets of possible worlds); and I have remained silent on how to analyze *might*-conditionals if CEM holds (roughly, as conditionals with modalized consequents). But the perhaps most pressing technical issue is that, as things stand, validity$_2$ not only delivers CEM and SDA but also the unwelcome principle of ANTECEDENT STRENGTHENING:

$$\text{(AS)} \qquad \phi > \chi \vDash (\phi \wedge \psi) > \chi$$

Lewis (1973) rejects AS since *Sobel sequences* appear to be perfectly consistent:

(7) If Alice had come to the party, it would have been fun. But if Alice and Bert had come to the party, it would not have been fun.

(7) is a sequence of counterfactuals but the point applies to conditionals of all stripes (see e.g. Willer 2017). And yet it is easy to see that no state that treats ϕ and ψ as possibilities can accept χ if strengthened with ϕ but reject χ if strengthened with ϕ and ψ.

Nothing about the antecedent strengthening facts undermines the important result that we can endorse CEM and SDA without IAT, but let me at least sketch how AS may be avoided in the current setting. The basic (and familiar) idea is that conditional antecedents may bring hitherto ignored possibilities into view and so while *ifs* are universal quantifiers over a set of possible worlds, this domain evolves dynamically as discourse proceeds (see von Fintel 2001 and Gillies 2007 for seminal discussion). So in entertaining the possibility of Alice coming to the party, we might ignore the possibility of Bert coming as well, and accept that the party will be fun on these grounds. Once we consider a conditional whose antecedent presupposes the possibility of Alice and Bill coming to the party, the modal horizon shifts, and it may very well be that all the possibilities thus brought into view are such that the party is not fun.

What underlies the invalidity of AS is then that one may entertain the possibility of ϕ without entertaining the possibility of $\ulcorner \phi \wedge \psi \urcorner$. Importantly, this response to the trouble with antecedent strengthening in a strict setting can be elaborated in such a way that we *preserve* the validity of SDA, the key observation concerning what it takes to entertain a disjunctive possibility (see Willer 2018 for details). One may, for sure, entertain the possibility of Alice coming to the party without Alice and Bert coming to the party. But in entertaining the possibility of Alice or Bert coming to the party, one is entertaining the possibility of Alice coming to the party, and the one of Bert coming to the party. The idea, in brief, then is that a conditional such as (8a) already brings into view the possibilities that matter for the antecedents of (8b) and (8c).

(8) a. If Alice or Bill come to the party, it will be fun.
b. If Alice comes to the party, it will be fun.
c. If Bill comes to the party, it will be fun.

The basic observation is that disjunctive possibilities imply each of their disjuncts as possibilities (Kamp 1973). So if conditionals presuppose the possibility of their antecedent, we expect them to simplify: the possibilities of the simplified conditionals are already in view and so there is no horizon shift. This marks the crucial difference with AS, where a horizon shift is very well possible. Clearly, this explanatory strategy does not effect the story about CEM, since $\ulcorner \neg(\phi > \psi) \urcorner$ and $\ulcorner \neg\phi > \psi \urcorner$ share the same possibility presuppositions.

I have not said much about the conceptual foundations of a bilateral approach to semantic theorizing. It is worth repeating, however, that the story about negation allows us to derive acceptance and rejection as incompatible attitudes, rather than having to stipulate that they are. As such, the theory faces no Frege-Geach style problems with negation. We thus have every reason to be confident that the bilateral semantics outlined here provides a conceptually and technically stable foundation for exploring the interplay between negation and conditionals in natural language.

Acknowledgments Warm thanks to three anonymous reviewers for their comments.

References

Maria Aloni. FC disjunction in state-based semantics, ms. Manuscript, University of Amsterdam.

David I. Beaver. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, Stanford, 2001.

Fabrizio Cariani and Simon Goldstein. Conditional heresies. *Philosophy and Phenomenological Research*, Forthcoming. http://dx.doi.org/10.1111/phpr.12565.

Kit Fine. Truthmaker semantics. In Bob Hale, Crispin Wright, and Alexander Miller, editors, *A Companion to the Philosophy of Language*, pages 556–577. Wiley-Blackwell, Malden, MA, 2nd edition, 2017.

Kai von Fintel. Bare plurals, bare conditionals, and *Only*. *Journal of Semantics*, 14(1):1–56, 1997. http://dx.doi.org/10.1093/jos/14.1.1.

Kai von Fintel. Counterfactuals in a dynamic context. In Michael Kenstowicz, editor, *Ken Hale: A Life in Language*, pages 123–152. MIT Press, Cambridge, MA, 2001.

Kai von Fintel and Sabine Iatridou. If and when *If*-clauses can restrict quantifiers. Manuscript, MIT, 2002.

Nissim Francez. Bilateral relevant logic. *Review of Symbolic Logic*, 7(2):250–272, 2014. http://dx.doi.org/10.1017/S1755020314000082.

Allan Gibbard. *Meaning and Normativity*. Oxford University Press, New York, 2013. http://dx.doi.org/10.1093/acprof:oso/9780199646074.001.0001.

Anthony S. Gillies. Counterfactual scorekeeping. *Linguistics and Philosophy*, 30(3):329–360, 2007. http://dx.doi.org/10.1007/s10988-007-9018-6.

James Higginbotham. Conditionals and compositionality. *Philosophical Perspectives*, 17(1):181–194, 2003. http://dx.doi.org/10.1111/j.1520-8583.2003.00008.x.

Luca Incurvati and Julian J. Schlöder. Inferential expressivism and the negation problem. In Russ Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 16. Oxford University Press, New York, Forthcoming.

Hans Kamp. Free choice permission. *Proceedings of the Aristotelian Society*, 74:57–74, 1973. URL http://www.jstor.org/stable/4544849.

Angelika Kratzer. Conditionals. *Chicago Linguistics Society*, 22:1–15, 1986.

David K. Lewis. *Counterfactuals*. Harvard University Press, Cambridge, MA, 1973.

Frank P. Ramsey. *The Foundations of Mathematics*. Routledge and Kegan Paul, London, 1931.

Iam Rumfitt. “Yes” and “No”. *Mind*, 109(436):781–823, 2000. http://dx.doi.org/10.1093/mind/109.436.781.

Paolo Santorio. Conditional excluded middle in informational semantics. In Alexandre Cremers, Thom van Gessel, and Floris Roelofsen, editors, *Proceedings of the 21st Amsterdam Colloquium*, pages 385–394, Amsterdam, 2017. ILLC Publications.

Mark A. Schroeder. *Being For: Evaluating The Semantic Program of Expressivism*. Oxford University Press, New York, 2008. http://dx.doi.org/10.1093/acprof:oso/9780199534654.001.0001.

Mark A. Schroeder. *Noncognitivism in Ethics*. Routledge, New York, 2010.

Robert C. Stalnaker. A defense of conditional excluded middle. In William L. Harper, Robert C. Stalnaker, and Glenn Pearce, editors, *Ifs: Conditionals, Belief, Decision, Chance, and Time*, pages 87–104. Reidel, Dordrecht, 1981. http://dx.doi.org/10.1007/978-94-009-9117-0_4.

Frank Veltman. *Logics for Conditionals*. PhD thesis, University of Amsterdam, 1985.

Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261, 1996. http://dx.doi.org/10.1007/BF00248150.

Malte Willer. Lessons from sobel sequences. *Semantics and Pragmatics*, 10(4):1–57, 2017. http://dx.doi.org/10.3765/sp.10.4.

Malte Willer. Simplifying with free choice. *Topoi*, 37(3):379–392, 2018. http://dx.doi.org/10.1007/s11245-016-9437-5.