Conditional Excluded Middle in Expressivist Semantics

Paolo Santorio

March 25, 2018

1 Introduction

This paper develops an expressivist theory of conditionals, with a broader view on a general expressivist theory of epistemic modality. Expressivism about conditionals, construed as the claim that conditionals don't express standard propositions, has a long pedigree. Philosophers of logic, building on Ernest Adams' foundational work (1975), have defended non-truth-conditional theories with the goal of vindicating a connection between conditional probabilities and probabilities of conditionals. Philosophers of language and formal semanticists have developed broadly expressivist theories of conditionals, mainly motivated by some logical properties of conditionals in natural language. (Veltman 1985, 1996; Yalcin 2007, 2012; Gillies 2004, 2009, among others). The view that I defend is indebted to both these traditions, but it also departs from them in two important ways.

The first difference concerns the motivation for rejecting a classical truth-conditional approach, which differs from standard arguments for expressivism. I start from a classical puzzle in conditional logic: theories of conditionals are faced with a tension between two plausible principles. On the one hand, conditionals seem to satisfy Conditional Excluded Middle, i.e. the principle that sentences of the form $(\phi > \psi) \vee (\phi > \neg\psi)$ are valid. On the other, conditionals of the form $\phi > \neg\psi$ seem incompatible with *might*-conditionals of the form $\phi > \Diamond\psi$. Unfortunately, these requirements are jointly unsatisfiable on a classical notion of consequence. I show that this tension is a close relative of a classical puzzle motivating expressivism about epistemic modality, and can be dissolved by adopting an expressivist framework for conditionals.

The second difference concerns the general view of the content and information expressed by conditionals. Standard expressivism about epistemic modality postulates a sharp divide between factual and epistemic information, roughly following the dichotomy between factual and normative content assumed by Gibbard (1990, 2003). On the view I defend, conditionals express both factual and nonfactual information. In other words, the assertion of a conditional can rule out factual possibilities (modeled as possible worlds); at the same time, conditionals also make distinctions along other dimensions. This mixed model throws an unexpected bridge between expressivism, on the one hand, and trivalent view of conditionals, on the other. Sometimes, factual information is sufficient to fix the value of a conditionals. In these cases, conditionals are straightforwardly true or false (at least, to the extent that standard declarative sentences are true or false).

The formalism backing this view builds on semantics that use information states (going back in particular to Veltman's 1985 data semantics). But just conditionals are treated differently than on standard informational semantics. The basic objects in the semantics are something like trajectories of evolution of information states—what I call 'information paths'. Intuitively, an information path represents a way that a subject's information might evolve, from its current state to a state of full certainty. The logic generated by this semantics is similar to the logic of informational/dynamic systems, but crucially it allows us to capture both Conditional Excluded Middle and the interaction between ordinary conditionals and *might*-conditionals.

I proceed as follows. In §2, I introduce the basic puzzle. In §3, I discuss and discard solutions based on homogeneity. §§4-5 give an intuitive presentation of the view, and §6-7 introduce a formal semantics and a notion of consequence. Throughout the paper, I stick to epistemic conditionals, but both the puzzle and (with substantial adaptations) the solution also extend to counterfactuals (for discussion of the latter, see Santorio 2018a).

2 A puzzle about conditional logic

I start from a classical puzzle about conditional logic, dating back to the debate on conditionals between Stalnaker 1968 and Lewis 1973. Conditionals seem to be subject to two plausible, but jointly incompatible constraints.

2.1 Conditional Excluded Middle

The first principle at stake is Conditional Excluded Middle.

Conditional Excluded Middle. (CEM) $\vDash (\phi > \psi) \vee (\phi > \neg\psi)$

The literature includes an impressive battery of arguments in favor of CEM. Here I rehearse two.

Argument #1: scopelessness. Epistemic conditionals with no overt modal appear to be scopeless with respect to logical operators. Importing and exporting these operators inside and outside the consequent of a conditional makes no difference to truth conditions. For reasons of space, here I only discuss negation, but the evidence for scopelessness includes the interactions between conditionals and quantifiers (Higginbotham 1986, von Fintel & Iatridou 2002), the adverb *only* (von Fintel 1997), and comparative constructions (Korzukhin 2014).

Notice, first of all, that negation can be imported inside and outside the scope of a conditional without affecting truth conditions. The sentences in (1) are equivalent.

(1) a. It's not the case that, if Frida took the exam, she passed.

b. If Frida took the exam, she didn't pass.

Notice also that the phenomenon persists with items that lexicalize negation, like *doubt* ($\approx$ *believe not*) and *fail* ($\approx$ *not pass*).

(2) a. I doubt that, if Frida took the exam, she passed.

b. I believe that, if Frida took the exam, she failed.

The lack of semantic interaction with negative items is perfectly expected on a theory that vindicates CEM (assuming standard Excluded Middle in the background logic), but not on theories that treat conditionals as universal quantifiers.

Argument #2: probability. The second argument is that, *modulo* plausible assumptions, CEM is required to vindicate basic judgments about the probability of conditionals. This point is underappreciated in the literature. I discuss it in detail elsewhere (Santorio 2017), but I give an overview here.

The link between conditionals and probability is notoriously complex. On the one hand, both intuition and experimental results suggest that speakers judge that, at least for unembedded conditionals, the probability of a conditional equals the conditional probability of the consequent, given the antecedent.[1] This suggests a generalization:

The Thesis. $Pr(\phi > \psi) = Pr(\psi \mid \phi)$
(for all ϕ, ψ, and for all Pr modeling rational credence)

On the other, so-called triviality results show that it is impossible to vindicate the Thesis in full generality, at least if we stick to standard semantic frameworks.[2]

In the face of this predicament, some theorists (e.g. Rothschild 2013[3]) have pushed an interesting line of thought. Triviality results show that the intuitive link between conditionals and probability is hard to vindicate in full generality. But, without settling whether this link can eventually be captured, we might take a shortcut and give a piecemeal vindication of the Thesis. I.e., we may show that most run-of-the-mill unembedded conditionals conform to the Thesis.

This piecemeal vindication of the Thesis seems highly desirable. At the very least, if a semantic theory fails to achieve it, it incurs the burden of explaining away the evidence supporting the Thesis. But, *modulo* some plausible assumptions, even this piecemeal vindication requires CEM.

Let me state the claim precisely. Assume that, if a conditional $\phi > \psi$ conforms to the Thesis, then the corresponding conditional with a negated consequent $\phi > \neg\psi$ also

[1]For a survey of classical experimental literature, see Evans & Over 2004; for more recent results that suggest that the Thesis is restricted to unembedded conditionals, see Douven & Verbrugge 2013.

[2]The literature on triviality results was started by Lewis 1976; see Hájek & Hall 1994 for an overview of classical triviality results. See also Charlow 2016 and Russell & Hawthorne 2016 for more recent results that also cover contemporary theories of conditionals and modality.

[3]Rothschild defends this line for the case of indicatives; Moss 2013 has put forward similar arguments for counterfactuals. Rothschild mention that the formal connection between the Thesis and Independence has been first pointed out by Ellis 1978.

conforms to the Thesis. (To see that this is intuitive, just check your own judgments about cases.) We can prove that, whenever $\phi > \psi$ has a consistent antecedent and conforms to the Thesis, the corresponding instance of the following probabilistic version of CEM holds.[4]

Probabilistic CEM (ProbCEM). $Pr((\phi > \psi) \vee (\phi > \neg\psi)) = 1$

In short:

> **Fact.** For all clauses ϕ, ψ (with ϕ consistent), and probability function *Pr*, such that $Pr(\phi > \psi) = Pr(\psi \mid \phi)$, the relevant instance of ProbCEM holds: $Pr((\phi > \psi) \vee (\phi > \neg\psi)) = 1$

Informally: if you assign $\phi > \psi$ a probability identical to $Pr(\psi \mid \phi)$, you must also assign probability 1 to the relevant instance of CEM.

Strictly speaking, Fact doesn't force us to adopt a *semantics* that vindicates CEM. We might be able to vindicate relevant instances of ProbCEM on a Lewis-style semantics, by cherry-picking an appropriate closeness relation in every relevant context. But this would be *ad hoc* and would rule out a number of intuitive choices for closeness.[5] By contrast, a semantics that validates CEM is immediately able to vindicate Fact.

---

[4] I build on the proof in Stefánsson 2014, which appeals to the following uncontroversial principle of conditional logic:

Conditional noncontradiction. (CNC) $(\phi > \psi) \supset \neg(\phi > \neg\psi)$

The proof:

i.	$Pr(\psi \mid \phi) + Pr(\neg\psi \mid \phi) = 1$	(probability calculus)
ii.	$Pr(\phi > \psi) + Pr(\phi > \neg\psi) = 1$	(i, Thesis)
iii.	$Pr((\phi > \psi) \wedge (\phi > \neg\psi)) = 0$	(CNC)
iv.	$Pr((\phi > \psi) \vee (\phi > \neg\psi)) = 1$	(ii, iii, probability calculus)

[5] For a simple illustration, consider a scenario where a fair 6-sided die has been tossed and you have no information at all about the coin toss. Suppose that someone says:

(i) If the die landed even, it landed on six.

Intuitively, (i) has probability 1/3. But this judgment is hard to vindicate on a semantics that uses universal quantification. To see this, take the following model. Let $D = \{w_1, w_2, w_3, w_4, w_5, w_6\}$ (subscripts correspond in the obvious ways to outcomes of the toss), and assume a background probability distribution *Pr* that assigns 1/6 to all worlds in *D*.

Assume a comparative closeness relation $\prec$ with the following features:

a. Strong Centering: for all w, $\prec$ ranks w as closer to itself than any world ($\forall w' \neq w : w \prec_w w'$.)

b. At each world w, all worlds different from w are tied for closeness ($\forall w', w'' \neq w : w' \approx_w w''$)

(a) is inherited from both Stalnaker's and Lewis's semantics. (b) is a natural assumption on an epistemic interpretation of closeness, since the evidential situation in the scenario is entirely symmetric. On top of this, assume a standard version of universal semantics in the Lewis-Kratzer style (using Greek metavariables to stand both for sentences and the relevant propositions):

$\ulcorner \phi > \psi \urcorner$ is true at w iff $\forall w'$ s.t. $w' \in max_{prec_w}(\phi)$, $w' \in \psi$

On this semantics, (i) is true when evaluated at w_6 and false when evaluated at all other worlds in the model. Hence, given *Pr*, (i) is assigned probability 1/6, contrary to intuition. To be sure, we *can* find some gerrymandered closeness relations that makes it the case that (i) gets probability 1/3 in this scenario. But these closeness relations will have counterintuitive results elsewhere; moreover, they are otherwise completely unmotivated.

2.2 *If* and *might*

The second principle at stake states the incompatibility of $\phi > \neg\psi$ and $\phi > \Diamond\psi$.

If-Might Contradiction. (IMC) $\quad (\phi > \neg\psi) \wedge (\phi > \Diamond\psi) \vDash \bot$

The evidence for IMC is straightforward. Pairs of conditionals of the relevant forms generate inconsistencies and disagreements, both in categorical contexts and under supposition.[6]

(3) # If Maria passed, Frida didn't pass; but, even if Maria passed, it might be that Frida passed.

(4) A: If Maria passed, Frida didn't pass.
B: I disagree. Even if Maria passed, it might be that Frida passed.

(5) # Suppose that, if Maria passed, Frida didn't pass, and that, if Maria passed, it might be that Frida passed.

Notice that IMC should be kept distinct from the following:

Duality. $\quad \vDash (\phi > \Diamond\psi) \leftrightarrow (\neg(\phi > \neg\psi))$

Several classical frameworks (e.g., Lewis 1973 Kratzer 1981, Kratzer 2012) make IMC and Duality equivalent. But, as I show in §7, the two can come apart.

2.3 Collapse

Given a classical notion of consequence, CEM and IMC together entail the equivalence of $\phi > \psi$ and $\phi > \Diamond\psi$. The direction $\phi > \psi \vDash \phi > \Diamond\psi$ is uncontroversial; as for the other direction:

i.	$\phi > \Diamond\psi$	Assumption
ii.	$\quad \phi > \neg\psi$	Supposition for conditional proof
iii.	$\quad \phi > \neg\psi \wedge \phi > \Diamond\psi$	(i, ii, $\wedge$-Introduction)
iv.	$\quad \bot$	(iii, IMC)
v.	$\neg(\phi > \neg\psi)$	(ii-iv, Reductio)
vi.	$\phi > \psi$	(v, CEM, Disjunctive syllogism)

Of course, this result is unacceptable. In response, classical theories drop one of CEM and IMC. Famously, Stalnaker (1968) endorses CEM and rejects IMC, while most other theorists, ranging from Kratzer (1981, 2012) to Gillies (2004) reject CEM. Both solutions are empirically costly, as the discussion in this section suggests.

[6] Stalnaker (1981, 1984) suggests that we reanalyze *might*-conditionals as involving an epistemic modal scoping over a conditional. Hence, the logical form of ⌜If ϕ, might ψ⌝ is $\phi > \Diamond\psi$. For reasons of space, I won't discuss this proposal in detail. But let me notice that it is incompatible with the widely accepted view that *if*-clauses can work as semantic restrictors of modals of various kinds, (see Kratzer 1986, 2012 for a classical defense of this view).

3 Trivalence by homogeneity

3.1 Homogeneity theories of conditionals

A number of theorists have recognized the tension between CEM and IMC and have offered a solution based on a trivalent semantics (von Fintel 1997, Schlenker 2004, Križ 2015 among others). The main maneuver is to claim that conditionals exemplify a phenomenon that is pervasive through natural language, i.e. homogeneity. A paradigm example of bearers of homogeneity are plural definite descriptions, like *the girls*. For illustration, consider (6):

(6) The girls passed the exam.

(6) is normally judged true if all the girls passed the exam, and false if none of them did (as originally pointed out by Fodor 1970). In situations where some but not all of the relevant girls passed, the sentence gives rise to a hedged judgment. As Manuel Križ has put it recently, "the natural answer is neither *yes* nor *no*, but *well*" (2015: 6).

(7) A:The girls passed the exam.
B:#Yes /#no /Well, only some of them did.

This kind of gappy behavior can be observed also under negation. By negating a sentence involving a plural description, we obtain a sentence that is true just in case the original sentence is false. Hence some scenarios are properly characterized neither by a sentence nor its negation. For example, neither (6) nor its negation in (8) exhaust the range of possible scenarios concerning the girls' passing or failing the exam. Neither of them is true in case only some of the girls passed.

(8) It's not the case that the girls passed the exam.

This suggests a natural analogy between plural descriptions and conditionals. On the one hand, sentences involving descriptions vindicate a kind of excluded middle principle. Whenever *The* ϕ *is* ψ has a defined truth value, exactly one between *The* ϕ *is* ψ and *The* ϕ *is not* ψ is true. On the other, a sentence of the form *The* ϕ *is* ψ is incompatible with a sentence of the form *Some* ϕ *is not* ψ. These principles parallel CEM and IMC, and encourage a homogeneity-based analysis of conditionals. Analyses of this kind are defended by a number of writers; here I give an informal overview of Schlenker's theory, though nothing I say depends on this choice.[7]

Schlenker analyzes conditionals as modal descriptions of antecedent worlds. Informally, the truth conditions of a conditional are glossed as follows:[8]

$\ulcorner \phi > \psi \urcorner$ is true at w iff the closest ϕ-worlds to w are ψ-worlds

[7] Both von Fintel and Schlenker treat homogeneity as a presupposition. Recently, Križ (2015) has convincingly argued that homogeneity effects in natural language are not presuppositional. Here I follow Križ, though nothing I say depends on this.

[8] On this gloss, as well as on Schlenker's theory, both conditionals and plural descriptions are nonmonotonic. As Schlenker makes clear, this is a non-essential feature of the theory and may be dropped.

In addition, Schlenker assumes a definedness condition, stating that conditionals have a definite truth value just in case the worlds in the domain of quantification behave homogenously with respect to the consequent.

> **Homogeneity Requirement (HR)**
> $\phi > \psi$ is true or false at w only if: either all closest ϕ-worlds to w are ψ-worlds, or all closest ϕ-worlds to w are $\neg\psi$-worlds

As it happens for descriptions, HR allows us to vindicate a version of CEM. *Modulo* a plausible semantics for *might*, IMC is also vindicated while avoiding collapse.[9]

3.2 Against homogeneity

While they still retain a broadly truth-conditional picture of meaning, homogeneity accounts are nonclassical, in the following sense: they generate a trivalent semantics, which admits a third truth value option besides truth and falsity. One upshot is that homogeneity accounts give rise to a nonclassical notion of consequence (see e.g. Cariani & Goldstein for discussion). My own proposal will share both of these features. But homogeneity yields the wrong kind of trivalence. Both plural descriptions and conditionals are trivalent, but they are trivalent for different reasons. This can be seen by noticing two points of divergence.

Probability judgments. I argued above that probability judgments are one of the reasons for endorsing CEM. But homogeneity-based accounts cannot help vindicate these judgments. In fact, existing experimental evidence suggests that there is a striking disanalogy between probability judgments about conditionals and about homogeneity-bearing expressions, like plural definites. To start, consider the following example.[10]

(9) If Maria flipped the coin, the coin came up tails.

Suppose that you know that the coin is fair and have no evidence one way or the other about how it landed. Plausibly, flip-and-heads-worlds and flip-and-tails-worlds are tied for closeness in your epistemic state.[11] In this situation, it seems rationally permissible to assign probability .5 to (9). Yet, if we assume HR, (9) suffers from homogeneity failure and hence doesn't have a defined truth-value. Now, it is unclear exactly what credence one should assign to statements that display homogeneity failure. But it seems both irrational and unusual to assign to them positive intermediate credence.

[9]As for CEM: we need the assumption that negation has a Strong Kleene semantics, taking falsity into truth. Given this, Schlenker's semantics, supplemented with HR, vindicates a weaker principle than CEM proper: $(\phi > \psi) \vee (\phi > \neg\psi)$ is valid *whenever* $\phi > \psi$ *is not undefined*. This is a kind of Strawson-validity, in the terminology of von Fintel 1999. As for IMC, we can proceed as follows. Treat $\phi > \Diamond\psi$ as involving existential quantification over the closest ϕ-worlds. Then $\phi > \Diamond\psi$ is incompatible with $\phi > \neg\psi$, but the entailment to $\phi > \psi$ is blocked. To be sure, we have that, whenever $\phi > \Diamond\psi$ is true, $\phi > \neg\psi$ is not false; but it can be that $\phi > \Diamond\psi$ is true and $\phi > \neg\psi$ is undefined.

[10]See Cariani & Santorio 2018 for a similar argument about *will*-sentences.

[11]I assume that the specific criteria we adopt for epistemic closeness won't matter here. If you think they do, just switch to a different example.

For a comparison, consider again (6).

(6) The girls passed the exam.

Suppose that half of the girls in the relevant set passed, and half of the girls did not. Ask yourself what credence you would assign to (6) in this scenario. While answers may vary, it would be very unusual to get the answer '.5'. In fact, the whole question of assigning a probability to (6), once we are aware that the homogenity requirement fails, seems somewhat confused or beside the point.

The intuitive disanalogy about probability judgments for conditionals and for plural definites is supported by experimental evidence. Cremers, Križ and Chemla (2017) find that probability judgments pattern in very different ways for, on the one hand, standard instances of homogeneity failure and, on the other, conditionals in contexts where not all possible antecedent worlds verify the consequent. In particular, when assigning probabilities to conditionals, a substantial number of subjects appear to ignore the frequency of 'gappy' cases: i.e., they ignore the frequency of cases where the conditional antecedent is false.[12] This is exactly what we expect on a view on which the probabilities of conditionals conform to the Thesis. Conversely, virtually no subjects use an analogous rule for assigning probabilities to statements involving plural descriptions. This is evidence that conditionals are not 'gappy', at least not in the way that sentences involving plural descriptions are.

Homogeneity Projection. In recent important work on homogenity, Križ (2015; see also Križ & Chemla 2015) points out that homogeneity projects in predictable ways in complex sentences. I.e., we can predict when a complex sentence will inherit the undefinedness of an embedded clause that it contains. The second disanalogy that I point out is that conditionals behave differently in this respect.

Consider first (10):

(10) All my students passed their exams.

As Križ points out (and as experimental work in Križ & Chemla 2015 has confirmed), (10) is judged true if all students have passed all of their exams, false if at least one student has failed all of theirs, and indeterminate otherwise. This suggests that sentences involving homogeneity triggers are evaluated in a supervaluational fashion, in the following sense. We check all the ways of resolving the undefined cases. If they all agree on a truth value for the sentence, the sentence is assigned that truth value; otherwise, it is left undefined. (Among other things, this explains why the sentence is heard as false if even just one of my twenty students fails all her or his exams.)

[12] I should note that this is not what the *majority* of the subjects do. This is puzzling, since other experimental work has shown that, in simple contexts, the majority of subjects' judgments do conform to the Thesis: see e.g. Baratgin et al. 2010. One natural conjecture is that the particular experimental design used by Cremers, Križ and Chemla encourages this kind of result. In any case, the important point is that their result shows a very clear difference between conditionals and homogeneity triggers.

Now, the cases that are of interest to us involve homogeneity triggers under non-monotonic operators.[13] Consider the following sentences.

(11) Exactly two of my students failed their exams.

(12) Usually, Frida passes her exams.

(13) In most classes, Frida handed in her assignments.

Focus on (11), and consider the following scenarios.

> *Scenario 1.* I have ten students. One of them failed all his exams, eight passed all of theirs and the remaining one failed some and passed some.
>
> *Scenario 2.* I have ten students. One of them failed all his exams, three passed all of theirs and the remaining six failed some and passed some.

As Križ 2015 and Križ & Chemla 2015 point out, we seem to have different judgments for the two scenarios. In Scenario 1, (11) is heard as undefined. Conversely, in Scenario 2, it seems straighforwardly false.

The comparison between Scenarios 1 and 2 teaches us that the simple supervaluational idea requires a twist. We said that a sentence involving a homogeneity trigger is evaluated as true/false if all ways of resolving the undefined cases yield truth/falsity, and is undefined otherwise. The extra twist is a restriction on how we resolve the undefined cases: they all have to be resolved in the same way. This explains the difference in judgments between Scenarios 1 and 2. Consider the distribution of truth values we get for all the relevant individuals in the two scenarios:

$$
\begin{array}{lll}
x_1 & \Rightarrow & 1 \\
x_2 & \Rightarrow & 1 \\
x_3 & \Rightarrow & 1 \\
x_4 & \Rightarrow & 1 \\
x_5 & \Rightarrow & 1 \\
x_6 & \Rightarrow & 1 \\
x_7 & \Rightarrow & 1 \\
x_8 & \Rightarrow & 1 \\
x_9 & \Rightarrow & \# \\
x_{10} & \Rightarrow & 0
\end{array}
\qquad
\begin{array}{lll}
x_1 & \Rightarrow & 1 \\
x_2 & \Rightarrow & 1 \\
x_3 & \Rightarrow & 1 \\
x_4 & \Rightarrow & \# \\
x_5 & \Rightarrow & \# \\
x_6 & \Rightarrow & \# \\
x_7 & \Rightarrow & \# \\
x_8 & \Rightarrow & \# \\
x_9 & \Rightarrow & \# \\
x_{10} & \Rightarrow & 0
\end{array}
$$

Scenario 1 Scenario 2

In Scenario 1, there is a way of changing all the undefined values in a uniform way that yields the result that exactly two of my students passed (namely, flip the only '#' to '0'). In Scenario 2, there is not. The only uniform ways of flipping the undefined truth values are ones on which exactly 9 students get a '1', or exactly 3 students do. (And in

[13] This discussion builds and elaborates on the examples and the discussion in Križ 2015, chapter 5.

fact, *Exactly nine of my students passed their exams* and *Exactly three of my students passed their exams* are intuitively undefined in Scenario 2.)

I have focused on (11) because it closely mirrors one of Križ's core examples. But analogous points can be made for the projection of homogeneity under other nonmonotonic quantifiers, such as *usually* and *in most classes*.

Now, the key observation is that conditionals pattern differently. In particular, we see no evidence for the 'resolve undefined cases in the same way' requirement. Consider:

(14) In exactly two parties, Maria danced if Frida danced.

And now take the following two scenarios:

> *Scenario 1**. There were ten parties in the semester, and Frida danced in all of them. We know that Maria danced in one and did not dance in eight of the others. We're uncertain about whether Maria danced in the last one.
>
> *Scenario 2**. There were ten parties in the semester, and Frida danced in all of them. We know that Maria danced in one and didn't dance in three. We're uncertain about whether Maria danced in the remaining six.

There is no difference in the way that (14) is evaluated in the two scenarios. In both cases, it is not judged either true or false. In particular, differently from what happened with (11) in the second scenario, we are unwilling to judge (14) false in Scenario 2*. This asymmetry is unexpected on a view that assimilates conditionals to plural definites.

This suggests that the evaluation of (14) doesn't exploit the kind of supervaluational reasoning that is the mark of homogeneity. Both plural descriptions and conditionals display nonclassical behavior, but this nonclassicality is due to different reasons .

4 Informational semantics

I start by outlining an analogy between the puzzle in §2 and a puzzle about epistemic modals discussed by Yalcin (2007). Yalcin points out that $\neg\phi$ and $\Diamond\phi$ seem to be inconsistent by the lights of some plausible tests. In particular, their conjunctions are infelicitous both when asserted by themselves and in embeddings:

(15) # It's not raining and it might be raining.

(16) # Suppose that it's not raining and it might be raining.

(17) # If it's raining and it might not be raining...

From a logical point of view, Yalcin's puzzle is triggered by the tension between the principle suggested by (15)–(17), on the one hand, and the requirement that $\Diamond\phi$ should not be veridical, i.e. that it should not entail ϕ, on the other.

Epistemic Contradiction. $\neg\phi \wedge \Diamond\phi \vDash \bot$

Nonfactivity of Epistemic Modality. $\Diamond\phi \nvDash \phi$

Both principles seem plausible, yet on a classical notion of logical consequence they are inconsistent.[14] Yalcin's suggested solution is to move to a nonclassical semantics, which generates a nonclassical notion of consequence. The resulting framework, which builds extensively on Veltman's update semantics for epistemic modality (1996), can accommodate both principles.

Here I'm not interested in establishing whether a Veltman-style semantics for epistemic modals is correct.[15] Rather, I want to highlight an obvious analogy between Yalcin's puzzle and the classical puzzle emerging from the Stalnaker/Lewis debate. This will provide motivation for developing a new solution to the latter, which builds essentially on the resources of informational and update semantics.

Start by noticing that our puzzle is generated by the tension between three plausible, but classically inconsistent principles.

Conditional Excluded Middle. (CEM)	$\vDash (\phi > \psi) \vee (\phi > \neg\psi)$
If-Might Contradiction. (IMC)	$(\phi > \neg\psi) \wedge (\phi > \Diamond\psi) \vDash \bot$
Nonfactivity of Might-Conditionals. (NMC)	$\phi > \Diamond\psi \nvDash \phi > \psi$

On this characterization, the analogy with Yalcin's puzzle is glaring. IMC and NMC are conditional counterparts of, respectively, Yalcin's Epistemic Contradiction and Nonfactivity principles. In the nonconditional case, the two principles are sufficient to generate inconsistency. In the conditional case, we also need CEM to draw a crucial inference from $\neg(\phi > \psi)$ to $(\phi > \neg\psi)$. But, aside from the extra assumption, the situation seems parallel.

I suggest that we build on this analogy. The classical Stalnaker/Lewis debate about CEM and *might*-conditionals should be reframed. On the usual construal, that debate is about the quantificational force of conditionals; a classical notion of consequence is simply assumed by default. I suggest that the tension between CEM and the behavior of *might*-conditionals points to a basic tension for classical semantics and logics for conditionals. Within a classical framework, this tension can be overcome only at a substantial empirical price. This motivates exploring nonclassical options, in particular options that build on informational and expressivist frameworks.

Below, I outline an expressivist theory of conditionals that solves the puzzle. Let me highlight that this is only one of the options that become available once we open up to nonclassical solutions. I choose it because it is both remarkably conservative

[14]Proof:

i.	$\Diamond\phi$	Assumption
ii.	$\quad\neg\phi$	Supposition for conditional proof
iii.	$\quad\neg\phi \wedge \Diamond\phi$	(i, ii, $\wedge$-Introduction)
iv.	$\quad\bot$	(iii, Epistemic Contradiction)
v.	$\neg\neg\phi$	(ii-iv, Reductio)
vi.	ϕ	(v, propositional logic)

[15]A number of theorists have sketched alternative accounts. See, among very many, von Fintel & Gillies 2011; Silk 2016; Mandelkern 2017

with respect to both truth-conditional and existing expressivist systems, and very well-behaved. But other options (including fully dynamic systems) might work equally well.

5 Epistemic expressivism and trivalence

5.1 The split model of content

Modern expressivism about epistemic discourse is explicitly inspired by Gibbard's expressivism about normative discourse (1990, 2003), and in particular by Gibbard's strategy for modeling nonfactual information.[16] Building on possible worlds semantics, Gibbard proposes an enriched model of logical space, on which possibilities are represented as pairs of a possible world and (something like) a normative standard. For short, call this a *world-norm pair*.[17] Possible worlds model factual information, i.e. roughly the information that is expressed by ordinary descriptive sentences. Conversely, normative standards model normative information, i.e. roughly the information that is expressed by ethical or normative claims. Factual claims rule out world-norm pairs on the basis of their world component, but not on the basis of their norm component; conversely for normative claims.

The separation between the factual and the nonfactual dimension tracks a kind of independence between the two. For Gibbard, there are no logical or conceptual entailments between factual and normative information. The split model fully vindicates this independence. All factual claims are consistent with all nonfactual claims; accepting a factual claim doesn't force us to adopt ay nonfactual claim, and *vice versa*.[18]

Many prominent brands of expressivism about epistemic modality explicitly build on the split model. For example, Yalcin (2007; 2011) suggests that we model contents by using pairs of a possible world and an *information state* (roughly, the set of possibilities compatible with the beliefs of a subject). Factual claims operate on the possible worlds dimension of content. Epistemically modalized claims operate on the information state dimension. For a simple example, consider the following sentences:

(18) a. It's raining.
b. It might be raining

On a Yalcin-style semantics, both claims in (18) are evaluated at a pair of a world and an information state $\langle w, i\rangle$. But (18)-a only imposes constraints on the world member of the pair, while (18)-b only imposes constraints on the information state member.

[16] For modern expressivism, see, among many, Yalcin 2007, 2011, Moss 2015, Swanson 2016.

[17] In Gibbard 1990, normative standards are understood as *systems of norms*, i.e. fully developed normative systems that specify, for every action, whether it is mandated, permitted, or forbidden. In Gibbard 2003, normative standards are understood as hyperplans, i.e. fully specific plans of action.

[18] The independence constraint descends from the challenge posed by the Open Question argument: for discussion, see (among very many) Darwall et al. 1992. Notice that the absence of a logical and conceptual entailments between the factual and nonfactual domain is compatible with the existence of other kinds of entailments, including grounding claims and *a priori* nonlogical entailments.

It's raining is true at $\langle w, i \rangle$ iff it is raining at w

It might be raining is true at $\langle w, i \rangle$ iff it is compatible with i that it is raining.

5.2 Entangling the factual and the epistemic dimension

My proposal builds on existing expressivist theories, but I refrain from endorsing the split model. Similarly to Gibbard, I take possibilities (understood as the basic points at which we evaluate sentences for truth and falsity) as involving both a factual and a nonfactual component. But I use both the factual and the nonfactual components to evaluate conditionals. As a result, conditional information is a mixed kind of information, which draws distinctions along both the factual and the nonfactual dimension. One upshot (if we assume that at least some claims in the language are assigned truth values) is that the semantics of conditionals is trivalent. Some conditionals are true and some are false, in the very same sense in which factual claims are true or false. Some other conditionals have a third status, which is characterized via a broadly nonfactualist picture of epistemic content. This strategy is at odds with standard trivalent pictures, on which conditionals that are neither true nor false are 'gappy' in the same way as sentences displaying homogeneity failure. At the same time, it is also a departure from standard expressivism, on which claims pertaining to epistemic discourse are never truth-apt.

I describe in detail the formal framework in §6, but I give an informal overview here. Before doing so, let me introduce some terminology. So far I have talked about 'factual information' and 'nonfactual information'. Since conditionals cut across both dimensions, I will talk about 'factual' and 'conditional information'.

On many expressivist semantics, epistemically modalized sentences like *must* ϕ or *might* ϕ are evaluated as true or false relative to an information state. For example, *must* ϕ is true at i if all the possibilities in i verify ϕ, and *might* ϕ is true at i if at least some possibility in i verifies ϕ. I suggest that we consider not only static epistemic states, but also their possible trajectories of evolution. Consider the epistemic state of an ordinary subject. That epistemic state settles some factual sentences as true, some as false, and leaves others undecided. If the subject acquires more information, the information state will grow and some sentences that were initially undecided will be settled as true or false. Eventually, the subject will reach a fully decided epistemic state, which settles all sentences as true or false. Of course, from the current standpoint there are many different trajectories that an epistemic state might take. Specifying all these possible trajectories is another way of specifying your current information state.

I suggest that these complete trajectories work as the basic model-theoretic objects in a semantics for conditionals. I call these trajectories *information paths*. Formally, you can think of a path as a sequence of smaller and smaller sets of worlds, up to a singleton.

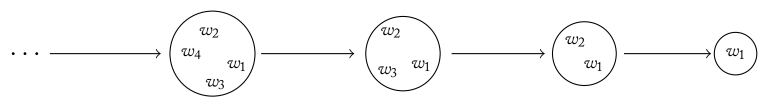

The endpoint of a path corresponds to the factual component in a Yalcin-style model of content. Hence it is obvious how we can use paths to evaluate purely factual sentences:

we just consider the endpoint of the path, disregarding the remaining structure. The latter structure is used to evaluate epistemic sentences.

So much for the basic model. Let me now explain informally how the semantics for conditionals works. The driving intuition is that assertions of conditionals are assertions performed under supposition. The effect of a supposition is modeled simply by removing from a path all the possibilities incompatible with the antecedent. After that, the consequent is evaluated as an ordinary clause at the updated path. Hence, for the case of bare conditionals (i.e. conditionals with no overt modal), the consequent is evaluated at the endpoint of a path, much like an ordinary declarative sentence. As an example, take (19):

(19) If Maria passed, Frida passed.

The antecedent *Maria passed* updates the path by removing all possibilities incompatible with it. At this point, the consequent *Frida passed* is evaluated at the updated path.

One upshot is that, when the antecedent of a bare conditional is true at the terminal world of a path, this world alone is sufficient to determine the truth value of the conditional. For example, if Maria passed, the truth value of (19) relative to a path depends exclusively on whether Frida passed. More in general, the factual component of a path allows us to specify sufficient (though not necessary) conditions for the truth of a conditional at that path.

5.3 Centering principles

We just observed that, in some cases, the terminal world of a path is sufficient to settle the truth value of a conditional at that path. In particular: if ϕ is true at the terminal world of a path, $\phi > \psi$ is true (false) at a path just in case ψ is also true (false) at the terminal node of the path. This semantic feature corresponds to well-known proof-theoretic principles in conditional logic, which are usually labeled 'Strong Centering' and 'Weak Centering'.[19]

Strong Centering. $\phi \wedge \psi \vDash \phi > \psi$
Weak Centering. $\phi > \psi \vDash \phi \supset \psi$ (for nonmodal and nonconditional ϕ, ψ)

Strong and Weak Centering are valid for epistemic conditionals in natural language. To see this, just run the usual tests: for example, notice that in both cases asserting the premise and denying the conclusion sounds contradictory. (The feeling of contradiction persists under embeddings of various sorts, including in suppositional contexts.)

(20) #Both Maria and Frida passed the exam. But it's not true that, if Maria passed, Frida passed.

(21) #If Maria passed, Frida passed. But Maria passed and Frida didn't.

[19] Weak Centering is often stated without restrictions to nonconditional and nonmodal ϕ, ψ. But the unrestricted version of Weak Centering leads to a well-known tension with the principle of Import-Export (see, among many, Gibbard 1981 and McGee 1985 for classical discussions, and Gillies 2009 and Khoo 2013 for recent ones). The weaker version sidesteps these difficulties.

Strong and Weak Centering also regulate probabilistic reasoning involving conditionals (on the assumption that the probabilities of simple conditionals conform to the Thesis).[20] The probability of (22)-a is the lower bound on the probability of (22)-b, exactly as a notion of probabilistic entailment would require (see Adams 1975).[21]

(22) a. Maria and Frida passed the exam.

b. If Maria passed, Frida passed.

Similarly, it is easy to prove that the probability of (23)-a is a lower bound on the probability of (23)-b.[22]

(23) a. If Maria passed, Frida passed.

b. Either Maria didn't pass, or Frida passed.

The vindication of Centering principles suggests that factual and conditional information are entangled. Given Strong and Weak Centering, conditionals stand in systematic relations of entailment with factual sentences. Moreover, these relations obtain merely on the basis of logical form, and hence are hardwired into the meaning of conditionals. This is an obvious difference with respect to factual and normative content: in that case, the lack of logical entailments between factual and normative information was exactly one of the factors motivating the split model. It is an advantage of the trivalent model that it vindicates Centering principles in such a natural way.[23]

5.4 Truth-aptness and trivalence

I have said that conditionals can sometimes be evaluated as true at a path just in virtue of the factual element in that path. By itself, this is not sufficient to say that conditionals are true or false *simpliciter*. Like the notions of truth at a world or truth at an

[20] For the observation that the notion of consequence that regulates probabilistic reasoning diverges from other notions of consequence that regulate epistemic discourse, see Schulz 2010 and Santorio 2018b.

[21] By definition, $Pr(A \mid B) = Pr(A \wedge B)/Pr(B)$. Since $0 \leq Pr(B) \leq 1$, it follows that $Pr(B > A) = Pr(A \mid B) \geq Pr(A \wedge B)$.

[22] Assume the Thesis, i.e. that $Pr(B > A) = Pr(A \mid B)$. Now, we have:

$$Pr(A \supset B) = Pr(A \wedge B) + Pr(A \wedge \neg B) + Pr(\neg A \wedge \neg B) =$$
$$Pr(A \mid B) \times Pr(B) + Pr(\neg B) =$$
$$Pr(A \mid B) \times (1 - Pr(\neg B)) + Pr(\neg B) =$$
$$Pr(A \mid B) - Pr(A \mid B) \times Pr(\neg B) + Pr(\neg B) =$$
$$Pr(A \mid B) + (1 - Pr(A \mid B)) \times Pr(\neg B)$$

Since $Pr(\neg B) \geq 0$, and since $Pr(A \mid B) \leq 1$, we know that $Pr(A \mid B)$ is a lower bound on $Pr(A \supset B)$.

[23] Of course, there are other ways of capturing Centering principles, some of which are available to split content theorists. In particular, given a simple informational semantics for conditionals in the style of Gillies 2009, Centering principles may be captured as informational entailments in the systems of Yalcin 2007 and Bledin 2015 (for the notion of informational consequence, see Yalcin 2007, Bledin 2015, as well as Veltman's notion of test-to-test consequence in his 1996). This may be a viable alternative. Let me notice though that Centering principles lack one of the marks of informational entailments: the latter are, in general, probabilistically invalid (see Schulz 2010, Santorio 2018b), while Centering principles are probabilistically valid.

information state, the notion of truth at a path is a technical notion that does work in the computation of semantic values for larger sentences. This notion does not, in general, correspond to the notions of truth an falsity that are applied to ordinary utterances.[24] However, we can supplement the recursive clauses that define truth at a path with a definition of truth *simpliciter*. For example, a natural way to do this is to adopt the following definition of truth:[25]

ϕ is true iff $[\![\phi]\!]^P = 1$, for all P where the the terminal element is $\{@\}$.

Informally, ϕ is true iff it is true at all paths terminating at the actual world. If we do this, we make conditionals truth-apt: they are the kind of statements to which (in some cases, at least) it makes sense to assign a truth value. Accordingly, on this interpretation the theory that I am defending is genuinely trivalent. This is a major divergence from standard epistemic expressivism, on which no claims pertaining to epistemic discourse count as truth-apt.

Nothing that I say in this paper depends on the full-fledged endorsement of trivalence. In particular, the resolution of the puzzle in §2 doesn't require introducing a notion of truth. But noticing that path semantics can be the compositional engine of a full-fledged trivalent theory is useful, for at least two reasons. First, it shows very clearly how the split model of content is abandoned. Second, it creates a new bridge epistemic expressivism with other trivalent accounts of epistemic discourse (see e.g. Mandelkern 2017). I don't attempt a comparison here, but it seems natural to think that this comparison (and perhaps attempts at integrating aspects of the two frameworks) could be theoretically fruitful.

6 Path semantics

6.1 Paths

The basic unit of path semantics is an *information path*. Informally, an information path is a sequence of information states that starts from the empty set and expands into a larger information state, adding one world at a time. Formally, we can define paths as follows. Let i be any information state, which I construe simply as a set of worlds. An information path (in i) is a sequence of subsets of i that is (i) ordered by subsethood and (ii) maximal, in the sense that there is no larger sequence of subsets of i that is ordered by subsethood. An information state uniquely determines a set of information paths; I use $\textsc{path}(i)$ to denote the set of paths determined by (or, as I will say sometimes, generated by) information state i.

It is useful to model paths via branching diagrams. For an example, here is the set of paths determined by the information state $i = \{w_1, w_2, w_3\}$. (To avoid clutter, I leave out the empty set, which is the beginning point of each path.)

[24] For relevant classical literature, see Dummett 1973 and Lewis 1980; see also Ninan 2010, Rabern 2012, Yalcin forthcoming, among others, for recent discussion.

[25] Notice that this is not the only way to do so. Notice also that (as it happens for supervaluational semantics) we will need to give a separate definition of falsity *simpliciter*.

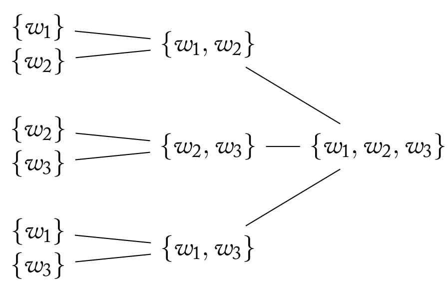

Intuitively, paths model possible ways in which information may grow. This becomes clearer by reading paths from right to left. A move to a smaller set in a path represents a possible transition from a less informed to a more informed state. (I will keep writing paths from the smallest to the largest set, left to right, because it makes the formalism more intuitive.)

6.2 Semantics

I state a semantics for a propositional language involving atomic sentences, Boolean connectives, epistemic modals, and conditionals. All sentences are evaluated relative to a path. This evaluation procedure is supplemented with a notion of update, which plays a key role in evaluating conditionals.

I take an **information path in** i to be a maximal sequence of elements of $\wp(i)$, ordered by the subset relation. I use the customary square brackets '$[\![\cdot]\!]$' notation for the interpretation function and relativize interpretation to a path parameter P. I also assume a background model $\langle W, V\rangle$, with W a set of worlds and V a valuation function mapping pairs of an atomic sentence and a world to $\{0,1\}$. Borrowing a formal device from Stalnaker 1968, I assume that W includes the absurd world λ, i.e. a world such that $V(\lambda, A) = 1$, for any A.[26]

These are the clauses for atomic sentences, connectives, and modals.

> Atoms: $[\![A]\!]^P = 1$ iff $w : \mathsf{min}(P) = \{w\}$, is s.t. $V(w, A) = 1$
> where $\mathsf{min}(P)$ is:
>
> - the smallest non-empty member of P, if P has a non-empty member;
> - the singleton $\{\lambda\}$, otherwise.
>
> $[\![\neg\phi]\!]^P = 1$ iff $[\![\phi]\!]^P = 0$
>
> $[\![\phi \vee \psi]\!]^P = 1$ iff $[\![\phi]\!]^P = 1$ or $[\![\psi]\!]^P = 1$
>
> $[\![\phi \wedge \psi]\!]^P = 1$ iff $[\![\phi]\!]^P = 1$ and $[\![\psi]\!]^P = 1$
>
> $[\![\Diamond\phi]\!]^P = 1$ iff for some $P' \in \textsc{path}(\bigcup P)$, $[\![\phi]\!]^{P'} = 1$

[26] The latter assumption just makes some features of the logic smoother. More realistic treatment of indicative conditionals will assume that conditionals include a presupposition to the effect that the antecedent is possible, following again an observation by Stalnaker (1975).

$[\![\Box\phi]\!]^P = 1$ iff for all $P' \in \textsc{path}(\bigcup P)$, $[\![\phi]\!]^{P'} = 1$

To give a semantics for conditionals, we first need to define the update $P[\phi]$ of a path P with a formula ϕ. To do this, we define two preliminary notions. The first:

ϕ is **true throughout** an information state i iff, for all P in $\textsc{path}(i)$, $[\![\phi]\!]^P = 1$

I.e.: ϕ is true throughout an information state i just in case it is true at all the paths that are generated by i. Second, we define the notion of the update of an information state with ϕ.

i' **is the update of** i **with respect to** ϕ (in short: $i[\phi]$) iff:

i. $i' \subseteq i$;
ii. ϕ is true throughout i';
iii. there is no larger set that meets conditions (i) and (ii).

In short: the update of i with respect to ϕ is the largest subset of i such that ϕ is true at all the paths generated by it.[27] It is easy to check that this yields intuitive results.[28]

Finally, we define the update of an information path P with respect to ϕ. This is just pointwise intersection of each information state in P with the updated information state that generates P:

Update of P with respect to ϕ: $P[\phi] = P \Cap (\bigcup P)[\phi]$

(with: $P \Cap i = \langle p_1 \cap i, \ldots, p_n \cap i \ldots \rangle$)

At this point, we can define truth at a path for conditionals in terms of update.

$[\![\text{if } \phi, \psi]\!]^P = 1$ iff $[\![\psi]\!]^{P[\phi]} = 1$

6.3 Examples

It is useful to see how a few sentences are evaluated at a sample path. For illustration, consider: ('M' and 'F' stand for the propositions that Maria passed and that Frida passed):

(24) $\langle \varnothing, \{w_{\bar{\text{M}}\bar{\text{F}}}\}, \{w_{\bar{\text{M}}\bar{\text{F}}}, w_{\text{M}\bar{\text{F}}}\}, \{w_{\bar{\text{M}}\bar{\text{F}}}, w_{\text{M}\bar{\text{F}}}, w_{\bar{\text{M}}\text{F}}\}, \{w_{\bar{\text{M}}\bar{\text{F}}}, w_{\text{M}\bar{\text{F}}}, w_{\bar{\text{M}}\text{F}}, w_{\text{MF}}\} \rangle$

Nonmodal sentences. Nonmodal sentences are invariably evaluated at the first nonempty set in a path, which by design contains a unique world. As a result, the semantics of nonmodal sentences is fully classical. Here are some examples of sentences that are true at (24).

(25)
a. Maria didn't pass.
b. Maria passed or Frida didn't pass.
c. Neither Maria nor Frida passed.

[27] Given the sentences we're able to express in the language, there will always be a unique such set.

[28] Some examples: for any nonmodal sentence ϕ, $i[\phi]$ is the set of ϕ-worlds in i; $i[\Diamond\phi]$ is i itself if i contains a ϕ-world, and $\varnothing$ otherwise; $i[\neg\phi \wedge \Diamond\phi]$ is invariably $\varnothing$. These predictions are in line with update semantics (Veltman 1996). An interesting, and in my view welcome, divergence: $i[\Box\phi]$ is identical to $i[\phi]$, i.e. i updated with $\Box\phi$ is the set of ϕ-worlds in i.

Modalized sentences. $\Diamond\phi$ and $\Box\phi$ are evaluated at the information state that generates the path: technically, this means that they are evaluated at the union set of the path. As a result, the semantics of modalized claims reduces to standard informational semantics. For illustration, (24) makes true:

(26)
a. It might be that Maria passed.
b. It might be that Frida didn't pass.
c. It might be that Maria didn't pass and Frida did.

Conditionals. Conditionals make full use of the path structure. Conditional antecedents update the path of evaluation; the consequent is evaluated at the updated path. Consider:

(27) If Maria passed, Frida passed.

We first use the antecedent *If Maria passed* to update the path. Starting from (24) and following the definition of update above, we get:

$$\begin{array}{ccccc} \langle\varnothing, & \{w_{\bar{\mathrm{M}}\bar{\mathrm{F}}}\}, & \{w_{\bar{\mathrm{M}}\bar{\mathrm{F}}}, w_{\mathrm{M}\bar{\mathrm{F}}}\}, & \{w_{\bar{\mathrm{M}}\bar{\mathrm{F}}}, w_{\mathrm{M}\bar{\mathrm{F}}}, w_{\bar{\mathrm{M}}\mathrm{F}}\}, & \{w_{\bar{\mathrm{M}}\bar{\mathrm{F}}}, w_{\mathrm{M}\bar{\mathrm{F}}}, w_{\bar{\mathrm{M}}\mathrm{F}}, w_{\mathrm{MF}}\}\rangle \\ \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\ \varnothing & \varnothing & \{w_{\mathrm{M}\bar{\mathrm{F}}}\} & \{w_{\mathrm{M}\bar{\mathrm{F}}}\} & \{w_{\mathrm{M}\bar{\mathrm{F}}}, w_{\mathrm{MF}}\} \end{array}$$

From here, removing redundancy:

(28) $\langle\varnothing, \{w_{\mathrm{M}\bar{\mathrm{F}}}\}, \{w_{\mathrm{M}\bar{\mathrm{F}}}, w_{\mathrm{MF}}\}\rangle$

At this point, we evaluate the consequent at the updated path. *Frida passed* is false at (28), hence the conditional is false at (24). Notice the key point that guarantees the validity of CEM: nonmodal consequents are always evaluated at a single world.

6.4 Analogies with existing frameworks: ordering semantics and data semantics

My presentation of the semantics crucially exploits paths, i.e. maximal sequences of information states. While this particular version of the formalism is somewhat new, it's important to point out that path semantics is quite conservative with respect to existing frameworks. In particular, paths can be 'read off' the model theory of other, more familiar semantics.

First, a path is equivalent to the pair $\langle i, \prec\rangle$ of an information state i and a total ordering $\prec$ on the worlds in i. Accordingly, we can think of the set $\textsc{path}(i)$ of the paths generated by i as specified by the pair $\langle i, S_{\prec}\rangle$, where $S_{\prec}$ is the set of all the total orderings on elements of i. Hence we can think of the models used in path semantics as the result of adding orderings to the models of standard informational semantics (Veltman 1996, Yalcin 2007). Second, paths appear in the models of Veltman's data semantics. Veltman (1985, p. 157) defines an *information model* as a triple $\langle S, \preceq, V\rangle$ such that: S is a non-empty set of information states; $\preceq$ is a partial ordering on S; and V

is a (partial) valuation function. The ordering $\preceq$ tracks the growth of information states: $s \preceq s'$ iff s contains at least as much information as s' (and possibly more). It's easy to see that information models already contain paths: in particular, paths are maximal chains in $\langle S, \preceq \rangle$.

Hence paths can be read off model-theoretic constructions that are familiar. At the same time, the way in which paths are used to evaluate conditionals is genuinely new. Both semantics based on orderings and Veltman's data semantics treat conditionals as a kind of universal quantifier (either over worlds, or over epistemic states). Conversely, in path semantics *if*-clauses are merely used to update the path of evaluation, and consequents are evaluated as independent clauses at the updated path. In a sense, from a strictly semantic point of view there are no conditionals at all in path semantics. There is no operator or connective that performs a specific semantic operation that conditionals express (much like in Kratzer's account for conditionals; see Kratzer 1986, 2012). Rather, conditionals are just ordinary clauses preceded by an updating device.

7 Logical consequence

7.1 Defining consequence

Path semantics allows us to define several notions of consequence. One captures preservation of truth at a path.

> **Path consequence.**
> $\phi_1, \ldots, \phi_n \vDash_P \psi$ iff for all paths P such that $[\![\phi_1]\!]^P = 1 \ldots, [\![\phi_n]\!]^P = 1$, $[\![\psi]\!]^P = 1$

While path consequence is useful for several purposes, it is not the notion that best captures what follows from a set of accepted premises. To see this, notice that path consequence fails to vindicate one of the signature inference patterns for reasoning of this kind, i.e. Łukasiewicz's principle.[29]

> **Łukasiewicz's principle.** $\neg\phi \vDash \neg\Diamond\phi$

But a notion of this sort can be easily defined. Recall first the definition of truth throughout:

> ϕ is **true throughout** an information state i iff, for all P in $\textsc{path}(i)$, $[\![\phi]\!]^P = 1$

Using this, we define:

> **Path-Informational consequence.**
> $\phi_1, ..., \phi_n \vDash_{PI} \psi$ iff, for all i such that $\phi_1, ..., \phi_n$ are true throughout i, ψ is true throughout i.

[29] See *[reference omitted for blind review]* for discussion of the links between reasoning and notions of consequence for natural language.

Path-Informational consequence is the analog, in the current framework, of Veltman's (1996) test-to-test validity, or Yalcin's (2007) informational consequence. Informally, it tracks what follows from an information state that validates certain premises. It is the obvious notion of consequence for assessing consistency and validity for asserted claims in natural language.[30]

7.2 Solving the puzzle

Path-Informational consequence vindicates both CEM and IMC, while blocking the collapse of *might*-conditionals onto bare conditionals (proofs are in the next section).

Fact 1. $\vDash_{PI} (\phi > \psi) \vee (\phi > \neg\psi)$

Fact 2. $(\phi > \neg\psi) \wedge (\phi > \Diamond\psi) \vDash_{PI} \bot$

Fact 3. $\phi > \Diamond\psi \nvDash_{PI} \phi > \psi$

Notice also that, despite the validity of IMC, Duality fails.

Duality. $\nvDash_{PI} (\phi > \Diamond\psi) \leftrightarrow (\neg(\phi > \neg\psi))$

This failure illustrates a difference between the logic that is generated by Lewis's classical semantics and my system. In Lewis's logic, IMC and Duality are equivalent. On Path-Informational consequence, Duality is stronger than IMC, hence it can still fail even though IMC is valid.

I claim that the validity of IMC and the invalidity of Duality is exactly what we see in the data. I already showed that IMC appears valid in §2. As for Duality: notice that the inference from (29) to (30) seems problematic.[31]

(29) If Maria passed, it might be that Frida passed.

(30) It's not the case that, if Maria passed, Frida didn't pass.

Hence, I claim, Path-Informational consequence does better than any of its classical counterparts at capturing the logic of conditionals in natural language.

7.3 Other logical features

It's useful to give a survey of other inference patterns that are validated by the two notions of consequence I have introduced.

Proposition 1. Restricted Import-Export is valid on Path Consequence.

[30] Besides Yalcin 2007, see Bledin 2015 for discussions of informational consequence. Informational consequence can be ultimately traced back to Stalnaker's (1975) notion of reasonable inference.

[31] To see this more clearly, consider (i) rather than (30) (which involves a kind of negation that is awkward in natural language).

(i) I doubt that, if Maria passed, Frida didn't pass.

(i) conveys that the speaker believes that, if Maria passed, Frida also passed. This is obviously too strong a conclusion to infer from (29).

Restricted Import-Export. $\phi > (\psi > \chi)_P \dashv\vDash_P (\phi \wedge \psi) > \chi$
(for nonmodal ϕ, ψ)

Proof. [To be added.]

Proposition 2. Might-Import-Export is valid on Path Consequence.

Might Import-Export. $\Diamond(\phi > \psi)_P \dashv\vDash_P \phi > \Diamond\psi$

Proof. Left-to-Right: take any path P such that $[\![\Diamond(\phi > \psi)]\!]^P = 1$. We have that, for some P' in $\textsc{path}(\bigcup P)$, $[\![\phi > \psi]\!]^{P'} = 1$, and hence $[\![\psi]\!]^{P'[\phi]} = 1$. Now, by *reductio*, suppose that $[\![\phi > \Diamond\psi]\!]^P = 0$. It follows that $[\![\Diamond\psi]\!]^{P[\phi]} = 0$, and hence that for all P'' in $\textsc{path}(\bigcup P[\phi])$, $[\![\psi]\!]^{P''} = 0$. But now, by the definition of path update, $P'[\phi] \in \textsc{path}(\bigcup P[\phi])$. Hence we also have that $[\![\psi]\!]^{P''} = 1$; contradiction. Right-to-left: by analogous reasoning, from $[\![\phi > \Diamond\psi]\!]^P = 1$ we know that, for some P' in $\textsc{path}(\bigcup P[\phi])$, $[\![\psi]\!]^{P'} = 1$. By *reductio*, suppose that $[\![\Diamond(\phi > \psi)]\!]^P = 0$. Hence, for all P'' in $\textsc{path}(\bigcup P)$, $[\![\phi > \psi]\!]^{P''} = 0$, and hence $[\![\psi]\!]^{P''[\phi]} = 0$. But $P''[\phi] \in \textsc{path}(\bigcup P[\phi])$; contradiction.

Proposition 3. Might-And-to-If is valid on Path Consequence.

Might-And-to-If. $\Diamond(\phi \wedge \psi)_P \dashv\vDash_P \phi > \Diamond\psi$

Proof. [To be added.]

Proposition 4. Łukasiewicz's principle is valid on Path-Informational Consequence.

Łukasiewicz's principle. $\neg\phi \vDash_{PI} \neg\Diamond\phi$

Proof. Suppose $\neg\phi$ is true throughout i, i.e. for all $P \in \textsc{path}(i)$, $[\![\neg\phi]\!]^P = 1$. By the clause for '$\neg$', we have that $[\![\phi]\!]^P = 0$. Hence, by the clause for '$\Diamond$', $[\![\Diamond\phi]\!]^P = 0$, hence $[\![\neg\Diamond\phi]\!]^P = 1$.

Proposition 5. Veridicality is valid on Path-Informational Consequence.

Veridicality. $\Box\phi \vDash_{PI} \phi$

Proof. Minor variant of the proof of Proposition 4.

Proposition 6. Or-to-If is valid on Path-Informational Consequence.[32]

Or-to-If. $\phi \vee \psi \vDash_{PI} \neg\phi > \psi$ (for nonmodal and nonconditional ϕ, ψ)

[32]It is unclear whether the restriction to nonmodal and nonconditional disjuncts is justified empirically. Intuitions are muddled by the fact that disjunctions of *might*-claims are subject to free choice effects (see Kamp 1973, Kratzer & Shimoyama 2002, among many). In any case, this restriction is not special to path semantics; rather, it is shared with standard informational/dynamic semantics for conditionals.

Proof. Suppose $\phi \vee \psi$ is true throughout i, i.e. for all $P \in \textsc{path}(i)$, $[\![\phi \vee \psi]\!]^P = 1$. Hence, for all $P \in \textsc{path}(i)$, either $[\![\phi]\!]^P = 1$ or $[\![\psi]\!]^P = 1$. Since ϕ and ψ are nonconditional and nonmodal, this means that, for all $w \in \bigcup P$, either $V(w, \phi) = 1$ or $V(w, \psi) = 1$. By *reductio*, suppose that $\neg\phi > \psi$ is not true throughout i, i.e. that for some P' in $\textsc{path}(i)$, $[\![\neg\phi > \psi]\!]^{P'} = 0$. Hence $[\![\psi]\!]^{P'[\neg\phi]} = 0$. Since ϕ and ψ are nonconditional and nonmodal, $V(w, \phi) = 0$ and $V(w, \psi) = 0$, with $\mathsf{min}(P') = \{w\}$. Contradiction.

8 Path semantics: a broader view

The project of path semantics goes beyond the puzzle discussed in this paper. Let me point to some further advantages and directions of research that I do not have the time to discuss here.

No covert modality. Standard truth-conditional semantics for conditionals posit the presence of a covert modal operator in all bare conditionals. For example, on a semantics like Kratzer's (1986, 2012), a conditional like (31)-a is analyzed as involving a covert necessity modal, as per the logical form in (31)-b.

(31) a. If Frida passed, Mary passed.
b. must [if Frida passed][Mary passed]

Vindicating CEM without appealing to homogeneity requires some rethinking. The relevant modal cannot be a necessity modal, as in (31)-b. We should rather postulate a covert selectional modal (i.e. a modal that 'selects' a unique world given a modal base), whose semantics is roughly along the lines of the selectional semantics of *will* in Cariani & Santorio 2018. This assumption is empirically costly. But, even assuming that natural language does contain some selectional modals, as Cariani and Santorio suggest, there is no clear example of an overt selectional modal with epistemic flavor. So we should assume that all natural languages contain a covert epistemic modal of this sort, and that this modal is in never overt. This seems implausible. Conversely, this assumption is not required on path semantics, which handles conditionals without appeal to any covert modality.

Generalization to counterfactuals. As I mentioned in the introduction, the tension between logical principles outlined in §2 is not confined to indicative conditionals. In fact, that tension first emerged just in the literature on counterfactuals. This suggests that path semantics should be generalized to other kinds of modality. I pursue this account in Santorio 2018a. Let me just notice that path semantics for counterfactuals involves nontrivial divergences in the semantics of *might*-conditionals and logical consequence. We need to depart more substantially from existing forms of expressivism.

9 Conclusion

Path semantics reconciles CEM and the inconsistency of $\phi > \neg\psi$ and $\phi > \Diamond\psi$, solving a problem that in various forms has been discussed since the beginning of modern work on conditionals. My proposal is confined to epistemic conditionals, but the puzzle generalizes to counterfactuals. The results of this paper encourage exploring the prospects for a general semantic framework for conditionals that accommodates both epistemic conditionals and counterfactuals.[33]

[33]Thanks to Maria Aloni, Fabrizio Cariani, Jennifer Carr, Paul Egré, Simon Goldstein, Daniel Greco, Manuel Križ, Benjamin Spector, Frank Veltman, and audiences at the ILLC, Institut Jean Nicod, the 21st Amsterdam Colloquium, and a philosophy of language workshop in Dubrovnik, Croatia.

References

Adams, Ernest Wilcox (1975). *The logic of conditionals: An application of probability to deductive logic*, vol. 86. Springer Science & Business Media.

Baratgin, Jean, David E. Over, and Guy Politzer (2010). "Betting on Conditionals." *Thinking and Reasoning*, 16(3): pp. 172–197.

Bledin, Justin (2015). "Modus Ponens Defended." *Journal of Philosophy*, 112(2): pp. 57–83.

Cariani, Fabrizio, and Simon Goldstein (????). "Conditional Heresies." Draft, Northwestern University and Lingnan University.

Cariani, Fabrizio, and Paolo Santorio (2018). "Will Done Better: Selection Semantics, Future Credence, and Indeterminacy." *Mind*, 127(505): pp. 129–165.

Charlow, Nate (2016). "Triviality For Restrictor Conditionals." *Noûs*, 50(3): pp. 533–564.

Cremers, Alexandre, Manuel Križ, and Emmanuel Chemla (2017). *Probability Judgments of Gappy Sentences*, Cham: Springer International Publishing, pp. 111–150.

Darwall, Stephen, Allan Gibbard, and Peter Railton (1992). "Toward Fin de Siècle Ethics: Some Trends." *Philosophical Review*, 101(1): pp. 115–189.

Douven, Igor, and Sara Verbrugge (2013). "The Probabilities of Conditionals Revisited." *Cognitive Science*, 37(4): pp. 711–730.

Dummett, Michael (1973). *Frege: Philosophy of Language*. Duckworth.

Ellis, Brian (1978). "A Unified Theory of Conditionals." *Journal of Philosophical Logic*, 7(1): pp. 107–124.

Evans, Jonathan St. B. T., and David E. Over (2004). *If*. Oxford: Oxford University Press.

von Fintel, Kai (1997). "Bare plurals, bare conditionals, and only." *Journal of Semantics*, 14(1): pp. 1–56.

von Fintel, Kai (1999). "NPI licensing, Strawson entailment, and context dependency." *Journal of Semantics*, 16(2): pp. 97–148.

von Fintel, Kai, and Anthony S. Gillies (2011). "'Might' Made Right." In A. Egan, and B. Weatherson (eds.) *Epistemic Modality*, Oxford University Press.

von Fintel, Kai, and Sabine Iatridou (2002). "If and When If -Clauses Can Restrict Quantifiers." Unpublished draft, available at http://web.mit.edu/fintel/fintel-iatridou-2002-ifwhen.pdf.

Fodor, Janet Dean (1970). *The linguistic description of opaque contexts*. Ph.D. thesis, Massachusetts Institute of Technology.

Gibbard, Allan (1981). "Two Recent Theories of Conditionals." In W. Harper, R. C. Stalnaker, and G. Pearce (eds.) *Ifs*, Reidel, pp. 211–247.

Gibbard, Allan (1990). *Wise choices, apt feelings*. Cambridge, MA: Harvard University.

Gibbard, Allan (2003). *Thinking How to Live*. Cambridge, MA: Harvard University.

Gillies, Anthony S. (2004). "Epistemic Conditionals and Conditional Epistemics." *Noûs*, 38(4): pp. 585–616.

Gillies, Anthony S. (2009). "On Truth-Conditions for If (but Not Quite Only If)." *Philosophical Review*, 118(3): pp. 325–349.

Hájek, Alan, and N. Hall (1994). "The Hypothesis of the Conditional Construal of Conditional Probability." In Ellery Eells, Brian Skyrms, and Ernest W. Adams (eds.) *Probability and Conditionals: Belief Revision and Rational Decision*, Cambridge University Press, p. 75.

Higginbotham, James (1986). "Linguistic Theory and Davidson's Program in Semantics." In E. LePore (ed.) *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Cambridge: Blackwell, pp. 29–48.

Kamp, Hans (1973). "Free choice permission." *Proceedings of the Aristotelian Society*, 74: pp. 57–74.

Khoo, Justin (2013). "A Note on Gibbard's Proof." *Philosophical Studies*, 166(1): pp. 153–164.

Korzukhin, Theodore (2014). "Dominance Conditionals and the Newcomb Problem." *Philosophers' Imprint*, 14(9).

Kratzer, Angelika (1981). "Partition and Revision: The Semantics of Counterfactuals." *Journal of Philosophical Logic*, 10(2): pp. 201–216.

Kratzer, Angelika (1986). "Conditionals." In *Chicago Linguistics Society: Papers from the Parasession on Pragmatics and Grammatical Theory*, vol. 22, pp. 1–15. University of Chicago, Chicago IL: Chicago Linguistic Society.

Kratzer, Angelika (2012). *Modals and Conditionals: New and Revised Perspectives*, vol. 36. Oxford University Press.

Kratzer, Angelika, and Junko Shimoyama (2002). "Indeterminate pronouns: The view from Japanese." In Y. Otsu (ed.) *Proceedings of the Tokyo conference on psycholinguistics*, vol. 3, pp. 1–25. Tokyo: Hituzi Syobo.

Križ, Manuel (2015). *Aspects of homogeneity in the semantics of natural language*. Ph.D. thesis, University of Vienna.

Križ, Manuel, and Emmanuel Chemla (2015). "Two methods to find truth-value gaps and their application to the projection problem of homogeneity." *Natural Language Semantics*, 23(3): pp. 205–248.
URL `https://doi.org/10.1007/s11050-015-9114-z`

Lewis, David (1976). "Probabilities of Conditionals and Conditional Probabilities." *Philosophical Review*, 85(3): pp. 297–315.

Lewis, David (1980). "A subjectivist's guide to objective chance." In *Ifs*, Springer, pp. 267–297.

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

Mandelkern, Matthew (2017). "Bounded Modality." Unpublished manuscript, Oxford University.

McGee, Vann (1985). "A Counterexample to Modus Ponens." *Journal of Philosophy*, 82(9): pp. 462–471.

Moss, Sarah (2013). "Subjunctive Credences and Semantic Humility." *Philosophy and Phenomenological Research*, 87(2): pp. 251–278.

Moss, Sarah (2015). "On the semantics and pragmatics of epistemic vocabulary." *Semantics and Pragmatics*, 8(5): pp. 1–81.

Ninan, Dilip (2010). "Semantics and the Objects of Assertion." *Linguistics and Philosophy*, 33(5): pp. 355–380.

Rabern, Brian (2012). "Against the Identification of Assertoric Content with Compositional Value." *Synthese*, 189(1): pp. 75–96.

Rothschild, Daniel (2013). "Do Indicative Conditionals Express Propositions?" *Noûs*, 47(1): pp. 49–68.

Russell, Jeffrey Sanford, and John Hawthorne (2016). "General Dynamic Triviality Theorems." *Philosophical Review*, 125(3): pp. 307–339.

Santorio, Paolo (2017). "Probabilities of Conditionals in Selection Semantics." Draft, University of California, San Diego.

Santorio, Paolo (2018a). "Counterfactuals without facts." Draft, University of California, San Diego.

Santorio, Paolo (2018b). "Credence and the Logic of Epistemic Discourse." Draft, University of California, San Diego.

Schlenker, Philippe (2004). "Conditionals as definite descriptions." *Research on language and computation*, 2(3): pp. 417–462.

Schulz, Moritz (2010). "Epistemic Modals and Informational Consequence." *Synthese*, 174(3): pp. 385–395.

Silk, Alex (2016). *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. Oxford: Oxford University Press.

Stalnaker, Robert (1968). "A Theory of Conditionals." In N. Recher (ed.) *Studies in Logical Theory*, Oxford.

Stalnaker, Robert (1975). "Indicative Conditionals." *Philosophia*, 5.

Stalnaker, Robert (1981). "A Defense of Conditional Excluded Middle." In W. Harper, R. C. Stalnaker, and G. Pearce (eds.) *Ifs*, Reidel, pp. 87–104.

Stalnaker, Robert (1984). *Inquiry*. Cambridge University Press.

Stefánsson, H. Orri (2014). *Decision Theory and Counter- Factual Evaluation*. Ph.D. thesis, London School of Economics.

Swanson, Eric (2016). "The Application of Constraint Semantics to the Language of Subjective Uncertainty." *Journal of Philosophical Logic*, 45(2): pp. 121–146.

Veltman, Frank (1985). *Logic for Conditionals*. Ph.D. thesis, University of Amsterdam.

Veltman, Frank (1996). "Defaults in Update Semantics." *Journal of Philosophical Logic*, 25(3): pp. 221–261.

Yalcin, Seth (2007). "Epistemic Modals." *Mind*, 116(464): pp. 983–1026.

Yalcin, Seth (2011). "Nonfactualism About Epistemic Modality." In A. Egan, and B. Weatherson (eds.) *Epistemic Modality*, Oxford University Press.

Yalcin, Seth (2012). "Bayesian Expressivism." *Proceedings of the Aristotelian Society*, 112(2pt2): pp. 123–160.

Yalcin, Seth (forthcoming). "Expressivism by Force." In D. Fogal, D. Harris, and M. Moss (eds.) *New Work on Speech Acts*, Oxford University Press.