Hedging your *if*s and vice versa

Kai von Fintel
MIT

Anthony S. Gillies
Rutgers University

The challenge

It's no secret that there are many competing views on the semantics of conditionals. One of the tools of the trade is that of any experimental scientist: put the object of study in various environments and see what happens. Edgington proposes a test:

> Any theory of conditionals has consequences for less-than-certain judgements. Something is proposed of the form: *If A, B* is true iff $A * B$. If a clear-headed person, free from confusions of a logical, linguistic or referential sort, can be nearly sure that $A * B$ yet far from sure that if A, B, or vice versa, then this is strong evidence against the proposal. (Edgington 1995: 260)

We will interpret this challenge as follows: explain how conditionals interact with expressions of less-than-certainty, including expressions of probability. This is a notorious problem and we will show that it is worse than thought.

To be clear from the outset: this paper is an advertisement for a problem. A good problem has a corresponding solution space that is interesting, with apparently easy-outs that prove to be deadends and sophisticated and tempting solutions that turn out to be insightfully wrong. What we will show is that this is, in fact, not a good problem but a great one.

1 The Compelling Intuition

Alex has views about the weather, the location of a certain marble, and whether the picnic will be worth going to. She might share some of this with Billy, staking herself to a flat-out claim:

(1) a. It's raining.
b. Blue is in the box.
c. The picnic won't be a success.

But perhaps her views are less-than-certain. Since the truthful speaker wants not to assert falsehoods she may instead issue a sentence with a hedge:

(2) a. It might be raining.
b. Blue is probably in the box.
c. It's unlikely that the picnic will be a success.

Hedging is a two-way street: since the trusting-but-truthseeking hearer wants not to agree to falsehoods, Billy may well deploy a hedge in (2) to resist what it is Alex asserts by going with one of the flat-out claims in (1).

There are (positive) hedges, the upshots of which weaken (assertions of) their prejacents. So a (positive) hedge on p is a modal expression plus p such that p asymmetrically entails the hedge on it. And then there are nearby negative hedges (usually incorporating negation), the upshots of which weaken denials of their prejacents. So a (negative) hedge like (2c) is a modal with an incorporated negation plus p such that the modal stripped of the negation plus $\neg p$ is a (positive) hedge on that.[1]

Since we are often ignorant and it's useful to communicate our relative ignorance, it's no surprise that we have conventional ways of doing that and no surprise that truthful speakers make use of those ways. What we have to say is in principle sayable about the gamut of hedges. But the problems we are interested in are (infamously) most pressing when the hedge is tied straightforwardly to probability. So we will stick (pretty much) to *probably* and nearby hedges here.

Hedges aren't the only way of expressing less-than-certainty. That's also the wheelhouse of ordinary indicative conditionals.

(3) a. He told Tom.
b. If he didn't tell Harry, he told Tom.

If Alex doesn't know that he (Billy) told Tom, she has little use for (3a) but she might have a use for (3b). So — no surprise — *if*s are hedges in their own right, weakening the claim made about the consequent.

A final non-surprise: hedges and conditionals intermingle with ease.[2]

1 Just so it's clear: we are interested in hedges that make a meaningwise difference to the sentences in which they occur. In particular: we will be interested in how *if* and *probably* interact because because the way you'd think they interact leads to triviality.

2 We'll be using (when it suits us) an intermediate propositional language L to model (part of) the relevant fragment of English we're interested in. We won't bore you with defining L properly, but it's equipped with some noteworthy operators: (i) two (epistemic) $\Box$s (one

(4) a. Probably if he didn't tell Harry, he told Tom.
$\textit{probably}\ (\textit{if}\ \ p)(q)$

b. Probably he told Tom, if he didn't tell Harry.
$(\textit{if}\ \ p)(\textit{probably}\ q)$

An ordinary *if* like (3b) can be hedged. That gives us (4a). As far as surface structure goes, this is a conditional in the scope of the hedge *probably*. And an ordinary prejacent like (3a) can be hedged and the result of that conditionalized. That gives us an iffed hedge like (4b). As far as surface structure goes, this is a conditional with the hedge *probably* embedded in the consequent. To have an umbrella term for these *hedge+if* constructions, call them conditional hedges.

Conditional hedges like those in (4) are sensible and useful ways of expressing relative ignorance. Strikingly, they seem to express the same thing. They both seem to express what we'll call the restricted meaning: they each are true iff the probability that he told Tom, given he didn't tell Harry, is high enough.

Compelling Intuition. Conditional hedges express restricted meanings:

$$\textit{probably}\ (\textit{if}\ \ p)(q) \ \Leftrightarrow\ \mathrm{Prob}(Q|P) \text{ is high enough } \ \Leftrightarrow\ (\textit{if}\ \ p)(\textit{probably}\ q)$$

This is still schematic, saying that conditional hedges and restricted meanings stand and fall together but not saying what that standing and falling is. Even in this to-be-filled-in form, the Compelling Intuition is compelling. But it also seems hard to square with what's possible: Lewis (1976) proved that the most natural route to filling it in and grounding it goes hopelessly off the rails. We want to explore how the Compelling Intuition can and — just as importantly — can't be saved. Part of that will mean stirring up more trouble for it.

Notice that we are not committed to the claim that the restricted meaning is the only possible meaning for the combination of a hedge and a conditional. In fact, we suspect that other readings are at least sometimes available. But

---

unary, one binary), (ii) two *probably*s (one unary, one binary), and (iii) a binary conditional connective $(\textit{if}\ \ \cdot)(\cdot)$. It will be clear why we double-up on modals and hedges (and context will disambiguate which is relevant). Not all the analyses we look at make use of all this, some of the paths diverging on just what expressive resources are best suited to modeling *if*s and hedges.

that is not our business here, which is the question of how to derive the restricted meaning when it arises as it so often does.[3]

Here is our plan. First we want to set-up the trouble as we see it for the Compelling Intuition (Section 2), and then sketch two ways out of the mess (Section 3). The basic strategy is to insist *if*s restrict the information over which the hedge is issued; the two ways are different implementations of that same basic strategy. But this is a good news/bad news situation. The Good: we have a robust diagnosis for the threat of triviality and a clear understanding of how to skirt it. The Bad: the two implementations fall prey to a more general threat of triviality — a threat we'll make vivid by what we call the Cross-Speaker Problem (Section 4). From here we scale a mountain of ever-better solutions (Sections 5–7). We won't speculate about whether where we end evokes (enough) summit-qualia. That's not our goal: we are here to chart the solution space thereby (re-)advertising the problem.

2 Trivialities, big and small

(The first part of this section is for aficionados. The executive summary: there is no sane way to give a semantics for conditionals and for hedges such that together, they result in the restricted meaning, thus justifying the Compelling Intuition. Hence, we need to look at some more or less insane ways.)

The Compelling Intuition needs to be filled in. The most natural way of filling it in says the $\Leftrightarrow$'s are the $\Leftrightarrow$'s of equivalence: that a conditional hedge is true (in a context, at a world) iff the relevant conditional probability is high enough.[4] Since conditional hedges are the result of layering a hedge like

3 There are unrestricted, doubly modal readings sometimes available. An old example we've used quite a bit:

(i) It is almost certainly false that if the die comes up even, it will be a six — since it almost certainly is a fair die.

This does not express that the conditional probability of six given given even is almost zero. To us these are interesting but (relatively) exotic readings, hence our preference to focus on conditional hedges that express the restricted meanings. Others have different preferences: the theory developed in Moss 2015, for instance, delivers exclusively layered, double-modalized meanings for hedges under *if*s (and so predicts that pairs like (4a) and (4b) can't be co-glossable).

4 Some additional conventions to keep the technical material as transparent as possible: we will take ordinary p's, q's, ... to stand in for arbitrary sentences, $[\![\cdot]\!]$ to be a (possibly partial) function from contexts and worlds to truth-values (and we'll treat $[\![p]\!]^c$ (sometimes omitting

probably with an *if*, we'd then have to give truth-conditions for each of them and derive (somehow) the desired equivalences. Let's focus on the first equivalence for now: *probably* $(if\ p)(q) \Leftrightarrow \mathrm{Prob}(Q|P)$ is high enough.

Assume that the hedge *probably* is a context-dependent modal, saying that (the proposition expressed by) its prejacent is high-enough:

Definition 1 (Probably)**.** $[\![\mathit{probably}\ p]\!]^{c,i} = 1$ iff $\mathrm{Prob}(P)$ is high enough where Prob is the (c,i)-relevant probability distribution.[5]

There are, of course, other options. But we are sticking with a simple truth-conditional framework for now and this simple picture makes the point we're making clear.

Rather than giving a semantics for *if* we'll put down two constraints on any way of doing things. We want to make the choices clear when we get into trouble. The first is a background constraint:

Constraint (Conditional Operator)**.** $[\![(if\ p)(q)]\!]^c = (C \cap P \to Q)$ where $\to$ is some conditional operator between $C \cap P$ and Q.

This just says that *if* has to express a relation between the propositions expressed by p and q.[6] The second constraint is The Equation. It is the first attempt we will see at an analysis that secures the Compelling Intuition:

Analysis 1 (The Equation)**.** $\mathrm{Prob}(P \to Q) = (\mathrm{Prob} + P)(Q)$, where $(\mathrm{Prob} + P)(\cdot)$ is the posterior probability got by updating $\mathrm{Prob}(\cdot)$ by P.[7]

This ties the probability of conditional propositions to the probabilities of the consequent proposition "updated" by the antecedent proposition. We haven't at this point said how the updating goes, just that the result of so updating must give us a proper posterior probability.

This package — that is The Equation together with the constraint that *if*s express propositions in virtue of expressing a conditional operator — seems

---

the 'c') as a function from worlds to truth-values (i.e., $[\![p]\!] \in 2^W$) giving the proposition expressed by p in c), and capital letters to stand in for the propositions expressed by their lower-case counterparts (thus $[\![p]\!] = P$). Since what possibilities an indicative cares about may well be context-sensitive, we'll also let C stand in for the c-relevant possibilities.

5 We'll require that Prob is (c,i)-relevant only if it is centered on C.

6 It doesn't say just which relation: see van Benthem (1986) (and then Veltman (1985) and Gillies (2010)) for a characterization result.

7 Notice that we haven't yet said what "updating" amounts to.

(were it not for looming triviality, about which more in a moment) to deliver (half of) the Compelling Intuition. Take any p and q:

$$\begin{aligned}[\![\textit{probably}\ (\textit{if}\ p)(q)]\!]^{c,i} = 1 &\text{ iff } \text{Prob}([\![(\textit{if}\ p)(q)]\!]) \text{ is high enough (Def. 1)}\\ &\text{ iff } \text{Prob}(C \cap P \to Q) \text{ is high enough (Constraint)}\\ &\text{ iff } (\text{Prob} + C \cap P)(Q) \text{ is high enough (Def. 1)}\\ &\text{ iff } (\text{Prob} + P)(Q) \text{ is high enough (Centeredness)}\\ &\text{ iff } \text{Prob}(Q|P) \text{ is high enough (Obviously?)}\end{aligned}$$

We start with a hedged *if* and we seem to end with what we want. It's the last step that is in trouble here: it assumes that the adding in the Equation — going from $\text{Prob}(\cdot)$ to $(\text{Prob} + P)(\cdot)$ — is conditionalization. But that's just what the Equation rules out. More precisely:

Fact 1. The Equation implies that $(\text{Prob} + P)(\cdot)$ is a (linear) mixture of $(\text{Prob}^1 + P)(\cdot)$ and $(\text{Prob}^2 + P)(\cdot)$ if $\text{Prob}(\cdot)$ is a mixture of $\text{Prob}^1(\cdot)$ and $\text{Prob}^2(\cdot)$.

The proof is straightforward: assume $\text{Prob}(\cdot)$ is a mixture, consider $\text{Prob}(P \to Q)$, and apply the Equation to $(\text{Prob}^1 + P)(\cdot)$ and $(\text{Prob}^2 + P)(\cdot)$. This is a strong constraint on the class of probability functions got to by "updating": it exactly characterizes imaging, and imaging and conditionalization only agree at the boundaries. Lewisian triviality then follows as a corollary.[8] So this way of grounding the Compelling Intuition won't work. And it's hard to see how else it could be grounded given this set-up.

2.1 NTV: Baby/bathwater

And so, some say, the set-up should go: Analysis 1 and it's companion Constraint won't do. The mistake was to think that the Compelling Intuition can be grounded by appeal to the truth-conditions of *probably* and *if*. Better to associate some other sort of property with indicatives and tie that property to conditional probability. There are different ways of trying to make good on this. One way that has proved popular: deny that ordinary bare indicatives traffic in ordinary truth-conditions at all and insist instead that they are devices for expressing (not reporting or representing) that the relevant con-

8 See Gärdenfors 1982.

ditional probabilities are high enough. This is the N(o)T(ruth)V(alue) camp.[9] Here is a broad and fairly neutral way of putting things.

Analysis 2 (NTV)**.** Pro-attitudes toward $(\textit{if}\ p)(q)$ (in c at i) co-vary with $\text{Prob}(Q|P)$ being high enough (where Prob is the (c,i)-relevant probability function). In particular: asserting $(\textit{if}\ p)(q)$ expresses that $\text{Prob}(Q|P)$ is high enough.

This doesn't say that $(\textit{if}\ p)(q)$ has as its conventional meaning that $\text{Prob}(Q|P)$ is high enough. But it does require — if this is to be a productive explanation — that the indicative conventionally encode the expression of conditional probability.[10]

This isn't wholly unattractive. Alex is still uncertain, but she is certain about this: either he told no one or he told one of $\{\text{Harry}, \text{Tom}\}$.

(5) a. If he told anyone, he told Tom.
b. If he told anyone, he told Harry.

Alex (let's say) doesn't want to sign-up for either of these, but she does want to distribute fine-grained pro-attitudes evenly between them. If the conditionals are mechanisms for expressing the relevant conditional probabilities this makes perfect sense: she signs up half-way for each conditional just because, given what she knows, the conditional probabilities are each $\frac{1}{2}$. This is progress.

But this, too, comes up short. That's because the traction that NTV views get by conventionally linking bare conditionals with the expression of conditional probabilities causes (immediate and sorta obvious) trouble when *if*s and hedges mix. Here it is the other half of the Compelling Intuition that makes the trouble stark: we get layered hedges instead of restricted meanings. Assume Alex has the conditionally appropriate pro-attitude toward a conditional hedge (rather than having a less-than-certainty attitude toward a bare conditional):

$$\begin{aligned}(\textit{if}\ p)(\textit{probably}\, q) \text{ iff } & \text{Prob}([\![\textit{probably}\, q]\!]|P) \text{ is high enough (Analysis 2)}\\ \text{iff } & \text{Prob}(\text{Prob}(Q) \text{ is high enough } |P) \text{ is high enough}\end{aligned}$$

9 Some campers: Adams (1975), Gibbard (1981), Edgington (2009), Bennett (2003), Bradley (2000).

10 There are different ways of implementing the basic idea, but those details won't matter for us.

Similarly if *probably* outscopes the conditional. Either way we end up with something a distance from, and a good deal less compelling, than the Compelling Intuition. The conditionals in (4) express the restricted meaning, not some weird and hard-to-parse layered restriction about what is probably probable.

The thing to say, from an NTV point of view, is that conditional hedges aren't just things that don't have truth-values but aren't even things that we take NTV-sanctioned pro-attitudes toward: when we think we are doing that toward $(\mathit{if}\ p)(\mathit{probably}\, q)$, we're really just expressing (not representing) that $\text{Prob}(Q|P)$ is high enough. But then we have completely given up on the idea that what we are up to in issuing conditional hedges can be understood by what we are up to in issuing *probably*s and *if*s. So this doesn't save or ground the Compelling Intuition; it throws it overboard. And so the Compelling Intuition is in some trouble.

2.2 Rothschild's Loophole

But wait. Maybe we were hasty in throwing out Analysis 1. We know it can't hold across the board, but maybe there is hope that it holds often enough in a well-behaved range of cases. Rothschild (2012) says the hope is well-founded. Broad brushstrokes: conditionals express conditional propositions all right (strict ones, in fact) but there are default constraints that, when in force, give us the Equation. And when those constraints aren't in force, we won't get the Equation and in particular won't be able to appeal its full-force in getting the triviality results.

A bit more detail: conditionals express strict conditional propositions over the possibilities compatible with X, an "idealized knowledge source" (knowing more than the speaker of the conditional, hedged or bare).

Analysis $1\frac{1}{2}$ (Centered Strict Conditional)**.**

i. Strict: $[\![(\mathit{if}\ p)(q)]\!]^{c,i} = 1$ iff $X \cap P \subseteq Q$ where X is the set of worlds compatible with the c-relevant idealized knowledge source.

ii. Constraints:

   a. Centering: If $[\![p]\!]^{c,i} = 1$ then $[\![(\mathit{if}\ p)(q)]\!]^{c,i} = 1$ iff $[\![q]\!]^{c,i} = 1$

   b. Independence: $\text{Prob}([\![(\mathit{if}\ p)(q)]\!] \cap [\![p]\!]) = \text{Prob}([\![(\mathit{if}\ p)(q)]\!]) \times \text{Prob}([\![p]\!])$

Centering amounts to requiring that the knowledge that X represents, given p, determines whether q; independence (a "default assumption") that the conditional is probabilistically independent of its antecedent. This package delivers: when the constraints are met, we get the Equation and thus that hedged *if*s — conditional hedges where the hedge scopes over the conditional — express restricted meanings.[11]

But there is trouble afoot. Centering doesn't play nice with conditional uncertainty. Sometimes our ignorance is structured in such a way that even conditional on p the status of q is wide open. And conditionals are a great way to communicate the truth about such situations. Suppose the worlds compatible with c are i, j, k with this distribution of truth-values:

	p	q	r
i	1	1	0
j	1	0	1
k	0	0	1

One way of expressing conditional uncertainty in a situation like this:

(6) That doesn't mean that if p, q. After all, maybe r — and if r, $\neg q$.

Residual uncertainty like this is at odds with centering.[12] An example: we don't know who did it, but we do know that of the three suspects only the driver must have acted alone. Pointing to the evidence collected thus far, one inspector is ready to declare that either it was the gardener or the butler. A voice of reason interrupts:

(7) That doesn't mean that if the gardener is innocent, the butler did it. After all, maybe it was the driver — and if it was, then the butler is innocent.

Signing up for (7) even if unbeknownst to the detectives the gardener is clean seems like signing up for a truth (and solid detective work).

11 What about conditional hedges where the hedge appears in the consequent of a conditional? For those Rothschild seems to go for a (somewhat less perspicuous) version of (a version of) Analysis 4. More on that shortly.

12 Two notes here: (i) the trouble we are raising is specifically about centering as a constraint on conditionals. What we have to say here is agnostic on, for instance, centering as a constraint on the covert necessity modal restricted by an *if*-clause (von Fintel 1997); (ii) conditionals resist widescope negation (as do some modals) but natural languages develop workarounds.

The same issue can be seen from a different angle. In this set-up it's important that X knows more than you do. That's because the truth of *probably* $(if\ p)(q)$ depends non-trivially on your (as issuer of the hedged *if*) credences — we want to allow that $\text{Prob}(\llbracket (if\ p)(q) \rrbracket)$ can take non-extremal values. Question: do non-probabilistic hedges like *perhaps/maybe/might* depend on X or on your more limited information? Suppose it's X. This would be weird since it predicts the consistency of things like

(8) There's a chance that it's raining but it can't be raining.

It's hard to hear that as anything but awful. So suppose the modals depend instead on your more limited information. That's no good either since instances of

(9) If p, maybe q and maybe $\neg q$.

can be true, useful, and informative things to say even if p is true.

(10) If the gardener's alibi checks out, then maybe it was the butler and maybe it wasn't. We'll have to keep investigating.

Barring some pretty extraordinary gymnastics, it's hard to see how to square this if the conditional here is strongly centered. (The gymnastics required aren't exactly what Stalnaker (1984) argues for in the case of (apparently) *might*-counterfactuals (that they are actually epistemic modals scoped over *would*-counterfactuals) but something more involved. The analog move for (10) is to say that it's not really a single conditional at all, but a conjunction of two conditionals each scoped under a *maybe*. It's harder to stick that landing.)

Maybe centering, like the independence constraint, comes and goes and when it goes we don't get the restricted meaning from a hedged *if*?[13] Nope. Alex and Billy watch Chris throw a die. Before they see the dots Alex utters (11a). Billy has options, (11b) and (11c) among them:

(11) a. Alex: If he didn't roll an even number, he rolled a 3.
b. Billy: Uh, no. It might have been a snake-eye.
c. Billy: That's unlikely.

13 We're exploring possibilities here. Rothschild takes the centering constraint to be non-give-upable.

In (11b) she could say something true even if unbeknownst to the two of them the roll was a 3. And what she would say by reaching for (11c) has the restricted meaning. So: it seems like we don't have strong centering (given her first option) and we have the restricted meaning (given her second).

All of that is to say: the Compelling Intuition is in some trouble.

3 Restricting Hedges

Somehow *if* s and hedges like *probably* conspire to express restricted meanings. As conspiracies go, this is relatively low-grade and easy to expose. In fact, it has been part of the folklore tradition in certain circles that this is how to respond to the challenge of Lewisian triviality.

Here is the basic idea. Assume that hedges make quantificational claims over a contextually relevant body of information — *probably* (in c at i) says that the proposition expressed by its prejacent has enough probabilistic mud on it (according to the (c, i)-relevant probability function). And suppose indicative *if*-clauses augment that body of information. Then if the augmenting takes precedence over the quantifying, we can save the Compelling Intuition from Triviality.

There are (at least) two ways of implementing that basic idea, and we'll sketch them here. The important points for us are that: (i) at their core they embrace this basic idea, (ii) the basic idea they implement (restricting over quantifying) is pretty awesome, and (iii) awesomeness notwithstanding, the basic idea faces trouble that these implementations can't handle as they stand.

Lewis (1975) argued that *if* s occurring under adverbs of quantification aren't conditionals at all. The stock example(s):

(12) $\left\{\begin{array}{r}\text{Always}\\ \text{Usually}\\ \text{Sometimes}\end{array}\right\}$ if a farmer owns a donkey he beats it.

On the surface — just like in our examples — these look like an *if* intermingling with an adverbial quantifier. They look like instances of (13a) or (13b):

(13) a. $\mathcal{Q}(\textit{if } p, q)$
b. $(\textit{if } p,\ \mathcal{Q}q)$
c. $\mathcal{Q}(p)(q)$

These too — just like our examples — express restricted meanings. But what single conditional operator could *if* pick out so that, when combined with an upstairs quantifier as in (13a), we get those meanings? No obvious answers. And what if it's a downstairs quantifier as in (13b)? Again, no obvious answers. So the restricting job these *if*s do isn't one a conditional operator can (easily) do. Better, Lewis said, to say that what *if* contributes in these environments is not some conditional operator with some iffy meaning but that it serves to mark an argument place in a polyadic quantifier as in (13c): this is a quantifier saying that $\mathcal{Q}$-many of the p-cases are q-cases. That's how the restricted meanings get expressed.[14]

The first implementation is inspired by this Lewisian line, but says its not only right when there are adverbial quantifiers nearby: it's right across the board and so right in the hedging cases we care about.[15] So all *if*s are devices for restricting other operators.

Our earlier examples:

(4) a. Probably if he didn't tell Harry, he told Tom.
b. Probably he told Tom, if he didn't tell Harry.

Since we're treating *probably* as an operator, if we insist that *if* expresses a bona fide conditional operator we have to sort out what scopes over what. Both options lead to trouble. We can side-step it though if we take these to be complex (polyadic) hedges, the *if*-clauses providing the restrictors.[16]

14 This isn't quite a knock-down argument: as we'll see, there are conditionals that can deliver the restricting behavior.

15 Kratzer develops and defends the idea (it's sometimes called "Kratzer's Thesis") in, for instance, Kratzer (1981, 1991, 1986, 2012). The obligatory quote: "The history of the conditional is the story of a syntactic mistake. There is no two-place *if . . . then* connective in the logical forms of natural languages. *If*-clauses are devices for restricting the domains of various operators" (Kratzer 1986: p. 11).

16 There's a bit of a mystery here. Lewis made a lot of hay of his triviality results about the interaction of probability and conditionals. At the same time, his discoveries in "Adverbs of quantification" seem very relevant to this topic. In fact, in an aside in that paper, he does hint that conditional hedges should be analyzed in the restrictor way:

> The <u>if</u> of our restrictive if-clauses should not be regarded as a sentential connective. It has no meaning apart from the adverb it restricts. The <u>if</u> in <u>always if</u> . . . , . . . , <u>sometimes if</u> . . . , . . . , and the rest is on a par with the non-connective <u>and</u> in <u>between . . . and</u> . . . , with the non-connective <u>or</u> in <u>whether . . . or</u> . . . , or **with the non-connective <u>if</u> in <u>the probability that . . . if</u> . . . .** It serves merely to mark an argument-place in a polyadic construction. [our emphasis]

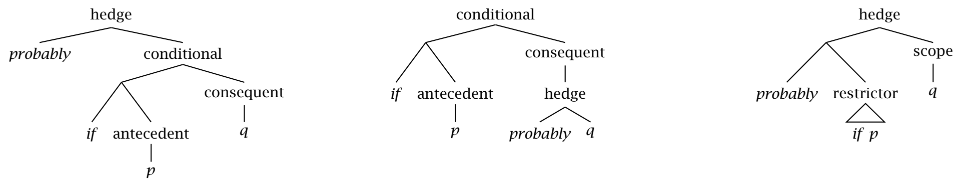

Figure 1 Hedge + *if*: Conditional vs. Restrictor

Analysis 3 (Restrictor)**.**

i. Tripartite: *hedge* + *if* constructions are polyadic hedges $\textit{probably}\,(\cdot)(\cdot)$.

ii. Restricting: $[\![\textit{probably}\,(p)(q)]\!]^{c,i} = 1$ iff $(\text{Prob} + P)(Q)$ is high enough where Prob is the (c, i)-relevant probability distribution.

The difference (i) makes can be put in pictures: Figure 1. The trees on the left feature either a (simple) hedge scoping over a conditional or a conditional with a hedged consequent; the one on the right features a (complex) hedge, the restrictor of which is given by the *if*.[17] There are lots of ways to spell-out a story about *probably* that meshes with all this. Since we don't much care how that goes, we illustrate in (ii) with the simplest picture.

Before, in the presence of the Equation, we couldn't take $(\text{Prob} + P)$ to go by conditioning Prob on P. Doing so forces extreme constraints on the class of credence functions. But now with Analysis 3 in place of the Equation, we don't face similar trouble. So we are free to insist that updating goes by conditionalization.

And doing that — insisting that $(\text{Prob} + P)(\cdot) = \text{Prob}(\cdot|P)$ — immediately delivers that conditional hedges express conditional readings. Conditional hedges like (4a) and (4b) have the same logical form: a hedge, restricted by whatever the restrictive material is in the *if*-clause, over the nuclear scope. Since *if*-clauses are devices for restricting operators, they augment the

On the other hand, it seems that Lewis did not anticipate that his ideas about restrictive if-clauses would have wider application. In Hajek's 1993 thesis, which Lewis supervised, there's a footnote that says: "David Lewis has suggested to me that 'if' should not always be analyzed as a sentential connective, but at least occasionally as a restriction modifier instead" (p.137, fn.72).

17 What goes for hedges goes for $\Box$, too. In particular, a bare conditional like *if it's not red, it's blue* is analyzed as a polyadic epistemic necessity claim $\Box(p)(q)$ with the *if*-clause contributing the restrictor.

contextually relevant information. The augmenting thus takes precedence over the quantifying, and we get the restricted meanings. Just as we wanted.

Analysis 3 requires re-thinking the structure of *if*s. Rather than going against the constraint that says that *if* expresses a conditional operator, we could instead re-think along dynamic lines how it expresses it. That's what the second implementation does and the upshot for us here is the same.

The most straightforward dynamic set-up takes semantic values not to be propositions (sets representing bodies of information) but context-change potentials (relations between bodies of information). Since hedges trade on what's likely, the bodies of information need to represent that, too. So we'll take the states of information to be pairs encoding the possibilities not yet ruled out (the non-gradable information we have in S) plus a representation of the plausibility information we have in S.[18]

Definition 2 (States)**.** An information state S is a pair $\langle C_S, \mathrm{Prob}_S \rangle$ where C_S is a set of worlds and Prob_S is a probability function centered on C_S. Two limit cases: **0** is any state in which $C_S = \emptyset$ and **1** any state in which $C_S = W$.

What p means (write it $[p]$) is a function from states to states — think of it as a set of instructions for changing what state we're in. The different sorts of information in a state have different instructions. Plain and simple sentences (no conditionals, no modals, no hedges) express the relatively plain and simple instructions of zeroing in on verifying worlds and updating the probabilistic information accordingly.

It is a bit easier (though not prettier to look at) if we first define an auxiliary kind of update $+$ and use that in the streamlined definition of $[\cdot]$. This way of doing things carves things where they should be carved: the obvious behavior of $+$ is separated out from the substantive hypothesis about *if*s and hedges interact in the dynamic set-up. So that's what we do here.[19]

18 We'll take that plausibility information to be tied to an underlying probability function, though in principle we could have opted for any number of semi-quantitative (conditional) plausibility measures — ranking functions being the closest to probability measures (Spohn 1988).

19 Officially we want to define $[\cdot]$ for the whole language and then turn to saying what it means for something to be true in a state. But it's just a little more convenient (and more accurately reflects what is up for grabs and what isn't) to bend the rules here a bit and define $[\cdot]$ for the *if*-free fragment and jump ahead to truth and then back to $(\textit{if}\ \cdot)(\cdot)$. So that's what we'll do. The officially story is that Definition 3, Definition 4 and Analysis 4 constitute a single recursive definition.

Definition 3. Let $S = \langle C_S, \text{Prob}_S \rangle$ be a state.

i. $C + p = C \cap P$ and $\text{Prob} + p = \text{Prob}(X|P)$ for atomic p

ii. $C + \neg p = C \setminus C + p$ and $\text{Prob} + \neg p = \text{Prob}(X|C \setminus C + p)$

iii. $C + (p \wedge q) = (C + p) + q$ and $\text{Prob} + (p \wedge q) = (\text{Prob} + p) + q$

Definition 4 (Simple Updates, Truth)**.** Let $S = \langle C, \text{Prob} \rangle$ be a state. Define $[\cdot]$ as follows (where L_0 is the *if*-/□-/hedge-free fragment of L):

i. $S[p] = \langle C + p, \text{Prob} + p \rangle$ for any $p \in L_0$

ii. $S[\textit{probably}\, p] = \begin{cases} S \text{ if } \text{Prob}_S(C_{S[p]}) \text{ is high enough} \\ \mathbf{0} \text{ otherwise} \end{cases}$

A sentence p is <u>true</u> in S, $S \models p$, iff $S[p] = S$.

Updating with plain non-hedged information winnows away at uncertainty and the probabilistic information gets updated as we go along. Hedges test whether the probabilistic information in the state is as the hedge says it is.

It's now easy to state the dynamic implementation of the restrict-before-quantifying idea.

Analysis 4 (Dynamic)**.**

i. Narrowscope: *hedge* + *if* constructions are narrowscoped hedges $(\textit{if}\ p)(\textit{probably}\, q)$.

ii. Conditional Updates: $S[(\textit{if}\ p)(q)] = \begin{cases} S & \text{if } S[p] \models q \\ \mathbf{0} & \text{otherwise} \end{cases}$

Two things to notice. First: as with Analysis 3 the first bit in (i) is a syntactic stipulation about the proper logical form for conditional hedges. Second: As we've set things up conditionals test to see if the conditional information is packaged in the state in the relevant way. The test can be put equivalently as checking whether $S[p]$ is a fixed-point of $[q]$.[20]

20 Modals work the same way: take □ to abbreviate $(\textit{if}\ \top)(\cdot)$. Analysis 3 and Analysis 4 both rule out hedges taking conditional connectives in their scope. This makes sense, especially since bare conditionals in these set-ups are themselves epistemically modal and hedges like *probably* do not naturally take a *must* in their scope. We've also assumed that hedges don't take other hedges in their scope.

This set-up also puts the restricting before the quantifying. Whether $(\textit{if}\ p)(\textit{probably}\ q)$ is true in S depends on whether *probably* q is true in the intermediate state $S[p]$. Since the truth of a hedge is context dependent (sensitive to your local state) this explains how *if*s and hedges conspire to express restricted meanings. Take any (plain) q and any state S. A simple inductive proof shows that $C_{S[q]} = C_S \cap Q$. Let $S = S[p]$. Then we have that $C_{S[p][q]} = C_S \cap (P \cap Q)$. And, of course, $\mathrm{Prob}_{S[p]} = \mathrm{Prob}_S(X|P)$. With that in mind:

$$
\begin{aligned}
S \models (\textit{if}\ p)(\textit{probably}\ q) &\text{ iff } S[p] \models \textit{probably}\ q \\
&\text{ iff } \mathrm{Prob}_{S[p]}(C_{S[p][q]}) \text{ is high enough} \\
&\text{ iff } \mathrm{Prob}_{S[p]}(C_S \cap P \cap Q) \text{ is high enough} \\
&\text{ iff } \mathrm{Prob}_S(Q|P) \text{ is high enough}
\end{aligned}
$$

Just as we wanted.

4 Cross-Speaker Problem

It's not all rosy, though. But before saying how the restrict-before-quantifying type of stories (in whatever implementation) don't actually ground the Compelling Intuition we want to mention and dismiss an objection to them. It may be all well and good, the objection grants, to appeal to syntactic wrangling (either by insisting that conditional hedges are really polyadic hedges or by insisting that the hedges featured in them are narrowscoped) when what needs explaining is why certain sentences's truth/acceptability/assertion-conditions/whatever covary with certain conditional probabilities. That's

---

We should also point out that what we've opted for here is a way of executing the dynamic program for conditionals that wears the dynamics on its sleeve (Gillies 2004). But there are other choices that make the connection to the standard set-up, and Analysis 3, more transparent (Gillies 2009, 2010). The alternative would still insist on narrowscoping the hedges but would add its own version of clause (ii.) that looks like this:

ii. Shifting: $[\![(\textit{if}\ p)(q)]\!]^{c,i} = 1$ iff $C \cap P \subseteq Q$ where $Q = [\![q]\!]^{c+p}$.

What, exactly, is $c + p$? It's the intermediate context got by taking c and adding the information p expresses to it. We're thinking here of contexts as providing two bits of information: a set of c-relevant worlds and a probability function centered on it. So $c + p$ is the result of updating each of those bits by P. For concreteness: take a context $c = \langle C, \mathrm{Prob}\rangle$. Then $c + p = \langle C \cap P, \mathrm{Prob}(X|P)\rangle$. We trust it's clear that the two dynamic stories come to the same thing. (Dynamic accounts of probabilistic hedges, like the one in Yalcin 2012 and the one here, follow Veltman's (1996) set-up for modals.)

fine if what needs explaining is some fact about language. But lurking just beneath the surface of the Compelling Intuition — arguably what gives it its compellingness — is that we can entertain certain propositions in thought and judge them likely and that those judgments seem to covary with the relevant conditional probabilities. (In fact, that's presumably the way that Edgington understands the challenge we quoted at the beginning.) Since thought isn't a place where syntactic wrangling happens, neither implementation of the restrict-before-quantifying idea gets at what needs getting at.

The objection gets traction (apart from the force it gets by the use of underlining (or, in variants, italics)) only if we can reliably tell the difference between (i) entertaining a proposition and judging it likely, and (ii) judging the truth of conditional hedges (even if unspoken). But that's the thing about language(-even-if-unspoken) and thought: it's hard to be sure judgments about one aren't coloring judgments about the other. We don't want to sign up for asserting or believing *All the beer in the fridge is cold* when there is no beer in the fridge. Is this non-signing up principally a property of what gets expressed by the bit of language giving voice to the thing we don't want to assert or believe? Or is our not wanting to say something giving voice to the thing we don't want to sign up for due to our grasping that that thing has some specific shortcoming? Like we said: if it's down to intuiting, it's hard to tell. Since the objection assumes that it isn't hard to tell, we think it doesn't carry much force against the restrict-before-quantifying sorts of strategy.

The problem with the restrict-before-quantifying strategy isn't that it's too closely tied to natural language.[21] The problem is that it's not tied close enough to other, nearby ways that hedges show up. Hedges make great rejoinders: when Alex utters one of (1) and Billy isn't so sure, the corresponding hedge from (2) make perfect sense. But Billy could also be economical in registering a hedge. For instance:

(14) a. Alex: He told Tom.
b. Billy: Probably so./That's likely.

Here the hedge scopes over an anaphoric device, which picks out Alex's claim.

So far, so good. But Billy could hedge in exactly the same way even if Alex had issued a conditional claim.

(15) a. Alex: If he didn't tell Harry, he told Tom.
b. Billy: Probably (so)./That's likely.

21 This argument has been floating around since at least von Fintel 2003.

Billy's hedge seems to operate on whatever the anaphoric devices here — *so* (possibly unpronounced) and *that* — pick out. And what Billy says expresses the restricted meaning: it's true iff the probability that he told Tom given he didn't tell Harry is high enough. But the anaphors seem to target Alex's conditional claim.

This is the cross-speaker problem.[22] In order for Lewisian triviality to threaten we need a conditional operator being scoped under a probabilistic hedge. Both implementations — Analysis 3 and Analysis 4 — get the hedges restricted without this. That was why they helped save the Compelling Intuition. But in (15) we seem forced to treat the hedge as scoping over something (the anaphor) that in turn has a conditional proposition as its semantic value. So we are back where we started.

We are not (yet) out of options, though. We'll survey what they are and how far they get us to where we want to get. We mention two options now so we can set them aside.

Despair Maybe we (speakers of natural language) are radically mistaken. The Compelling Intuition might be compelling but the connections it claims between conditional hedges and conditional probability do not obtain.

Copy/Reparse We can save the restrict-before-quantifying approach by saying that the anaphoric devices target linguistic strings (not their contents). The resulting hedge then needs reparsing, which delivers either a polyadic hedge (for Analysis 3) or a narrowscope hedge (for Analysis 4).

Take each of these in turn.

Despair doesn't really stand on its own: since there are various links in the Compelling Intuition wherein our error might lie. So it needs pairing with a thesis about whether or not conditional hedges have truth-conditions.[23]

22 It's poorly named. The trouble it exposes doesn't really rely on multiple speakers.

(i) Did he tell Tom if he didn't tell Harry? Probably.

The hedge here makes perfect sense, and has the restricted meaning, whether there is another speaker who asks the question or not. Nor does it matter whether (i) is spoken out loud at all. Dialogues like the ones we consider in the text do make it vivid, though.

23 Jackson (2006) pairs Despair with NTV: "Our usage of the indicative conditional construction is governed by a mistaken intuition [. . .]. We [. . .] wrongly think and speak as if the indicative

Plumping for a massive error theory isn't going to save anything. Maybe there are things to be said here, but they won't be reasons having to do with how close it gets us to saving the Compelling Intuition.

We understand Despair. We do not really understand a motivation for Copy/Reparse. Indeed we know of no one staking claim to this sad little patch of logical space. It seems dodgy and ad hoc and like the last place we'll find a satisfying way of grounding the Compelling Intuition. That's not a knock-down argument, but it doesn't need to be.

5 Divide and Conquer

There is a principled response to shield the restrict-before-quantify approach from the cross-speaker problem.[24] It starts by noticing a close cousin to hedging: (quantifier) scaling.

Alex thinks there is a problem in the department:

(16) Every student smokes before class.
$every\,(p)(q)$

Billy thinks she's overstating things. Some options:

(17) a. Most students smoke before class.
b. Most (of them) do/do that.
c. Most (do/do that).

No matter which of these Billy chooses, what Billy expresses is the same thing and it's a scaling back of Alex's quantificational claim. That means that there are two anaphors lurking. There's the VP anaphor (it picks up the nuclear scope for the quantifier). But since (17b) and (17c) both say the same thing as (17a), it also seems like *most* comes equipped with its own covert anaphor to pick up the restrictor.

What's good for quantifiers is good for modals. Since hedges are polyadic modals (assuming Analysis 3), we have our response to the cross-speaker problem. Here is our earlier example again:

conditional in fact has truth conditions such that its probability is the conditional probability of its consequent given its antecedent." It's open to go instead with a story on which hedged conditionals get truth-conditions but that we are in error in thinking they line up in an interesting way with conditional probabilities. This variety of Despair is broadly Lewisian in spirit.

24 Kratzer (2012) explores this strategy.

(15) a. Alex: If he didn't tell Harry, he told Tom.
b. Billy: Probably (so)./That's likely.

The explicit (even if unpronounced) anaphors *so* and *that* pick out the scope of the hedge. Hedges, being (semantically speaking) quantifiers, also have their homegrown covert anaphors. Call it α. That gives us a logical form for (15b) along these lines:

(18) *probably* (α) (*so*)

Since $[\![\textit{so}]\!] = [\![\textit{he told Tom}]\!]$ and $[\![\alpha]\!] = [\![\textit{he didn't tell Harry}]\!]$ Billy's reply is true iff the probability that he told Tom given he didn't tell Harry is high enough. That's the restricted meaning, and we steer clear of triviality. Time for a drink?

Not yet. The key to this strategy is that the two anaphors are independent and can roam for their targets free from one another. With scaling we can make use of an explicit (even if unpronounced) anaphor to get the nuclear scope. But using an anaphoric device ain't obligatory. We could go for the scope outright by just saying explicitly what it is *do* picks out:

(19) Most (of them) smoke before class.

This is a perfectly good — and equivalent — way to scale back Alex's universal claim in (16). It's here that hedging and scaling appear to part company. We can go for the scope of a polyadic hedge outright, too, but when we do sometimes we can't thereby express the restricted meaning. That is enough to derail this strategy.

We can make this vivid by comparing ways Alex can reply to Billy's different replies.

(20) a. Alex: If he didn't tell Harry, he told Tom.
b. Billy: Probably (so)./That's likely.
c. Alex: No really. I know he told one of them.

This is just our original example extended with Alex's reply in (20c). Billy's hedge has the restricted meaning, and Alex is resisting that hedge.

(21) a. Alex: If he didn't tell Harry, he told Tom.
b. Billy: Probably he told Tom./It's likely he told Tom.
c. Alex: ??No really. I know he told one of them.

In (21b) Billy doesn't rely on an anaphoric device, going outright for the scope of Alex's claim. But the resulting hedge doesn't — can't — thereby express the restricted meaning. That's why Alex can't reply by resisting the hedge. It's better for her to agree: that's why she went for the conditional in the first place. So whatever is going on in the cross-speaker problem, it is not that *probably* covertly gets its restriction from the original bare conditional. And so appeals to something like this will not help the restrict-before-quantify strategy to the problem posed by conditional hedges.

There's nothing special about *probably* in this respect. Nearby hedges pattern in the same way. A non-exhaustive catalog:

(22) a. That may be true./That's plausible./I guess that's true.
b. Maybe he told Tom./It's plausible he told Tom./I guess it's true he told Tom.

(23) a. I doubt that's true./That's unlikely./Seems fifty-fifty to me.
b. I doubt it's true he told Tom./It's unlikely he told Tom./Seems fifty-fifty to me that he told Tom to me.

In contexts like the one we've been considering, the anaphora-laced hedges in (22a) and (23a) can express the relevant restricted meanings but their explicit cousins in (22b) and (23b) can't. Thus the explanation for why *probably so* manages to express the restricted meaning cannot be that *probably* inherits the restrictor from the antecedent of the upstream conditional claim and *so* plugs in the value for the consequent: when we put that value in directly, we don't get the right meaning out.[25]

25 It has been put to us (independently by Angelika Kratzer and Justin Khoo) that the Divide & Conquer approach might be rescued. We are urged to note that the problematic cases where one doesn't get the restricted meaning are all cases where the overt complement of the hedge is a finite indicative past tense clause and that the cases where the restricted meaning emerges are all cases where the complement is an anaphor (*so*, *that*) and thus not (on the surface) a finite indicative past tense. Perhaps one could tell a story that the hedge can only anaphorically pick up the *if*-clause if its argument is in the right mood. We are skeptical that such a story could be told but of course we're willing to listen.

For now, we would like to see whether there are ways of getting the restricted meaning from a structure like *hedge* + α, where α is an anaphor picking up its antecedent from a prior conditional. If this is possible, we may have a new way of grounding the Compelling Intuition.

6 Belnapian Restriction

Lewis's argument that no conditional operator can do the restricting job of *if*-clauses under adverbs of quantification — and Kratzer's extension of it — isn't airtight. We've already seen one counterexample: whatever their other virtues and vices, dynamic implementations like Analysis 4 can give us a conditional operator that does the restricting that needs doing. Lewis knew of another: the Belnap alternative.[26]

Here is the idea. Conditionals express propositions, but gappy ones: if things are as the *if*-clause says, the conditional says what its consequent says; otherwise nothing. We also need to amend the basic story about modals and hedges. Since they can outscope conditionals, and conditionals might not get a truth-value at a point of evaluation, we need them to only care about points at which their prejacents get a proper truth-value. That makes sense.

Analysis 5 (Belnapian)**.**

i. Widescope: *hedge* + *if* constructions are widescoped hedges

ii. Gappiness: $[\![(\textit{if}\ p)(q)]\!]^{c,i} = [\![q]\!]^{c,i}$ if $[\![p]\!]^{c,i} = 1$ (undefined otherwise).

This pairs with the reigned-in quantificational oomph of modals and hedges:

Definition 5 (Restricted Quantification)**.** Let $\triangle$ be a modal with quantificational force $\mathcal{Q}$ and let Prob be the c-relevant probability function.

i. $[\![\triangle p]\!]^{c,i} = 1$ iff for $\mathcal{Q}$-many j's such that j is a c-relevant world and $[\![p]\!]^{c,j}$ is defined: $[\![p]\!]^{c,j} = 1$.

ii. $[\![\textit{probably}\ p]\!]^{c,i} = 1$ iff $\frac{\text{Prob}([\![p]\!])}{\text{Prob}(\{j\colon [\![p]\!]^{c,j} \text{ is defined}\})}$ is high enough.

For instance: if $\Box$ cares only about C then $[\![\Box p]\!]^{c,i} = 1$ just in case all of the worlds in C at which p gets a proper truth-value are worlds where that truth-value is 1. The hedging clause says that what *probably* cares about is the ratio of the probability that the prejacent is true to the probability that it has a proper truth-value. This reigning-in has no effect if the prejacent isn't

26 See Belnap 1970. Lewis dismisses this line for dealing with *if*s under adverbs of quantification, saying the price for going for such a "far-fetched" story about *if* is "exorbitant". The basic strategy is being resuscitated by Huitink (2009). She doesn't talk about hedges but the idea is clear. We think Lewis is right that, in the end, this strategy isn't worth the price, but it is worth a longer-than-a-footnote test-drive.

gappy. The gaps, introduced by the conditional, is what allows the restricting job to get done.

It's easy to see that secures the Compelling Intuition. First notice that $[\![(\textit{if}\ \ p)(q)]\!]^{c,j} = 1$ iff $j \in P$ and $j \in Q$. So $[\![\textit{probably}\ (\textit{if}\ \ p)(q)]\!]^{c,i} = 1$ iff

$$\frac{\text{Prob}(P \cap Q)}{\text{Prob}(\{j\colon\ [\![(\textit{if}\ \ p)(q)]\!]^{c,j} \text{ is defined}\})}$$

is high enough. But since $(\textit{if}\ \ p)(q)$ only gets a truth-value at worlds where p is true:

$$\text{Prob}(\left\{j\colon\ [\![\textit{if}\ p, q]\!]^{c,j} \text{ is defined}\right\}) = \text{Prob}(P)$$

So the hedged *if* is true iff

$$\frac{\text{Prob}(P \cap Q)}{\text{Prob}(P)}$$

is high enough. Since that is just the ratio formula for $\text{Prob}(Q|P)$ we have exactly what we want. And there is no cross-speaker problem: there is a conditional operator in Alex's claim for Billy's anaphor to pick out, that is what the hedge has scope over, and the resulting meaning is the correct restricted meaning.

But there is trouble. The trouble isn't so much that the required gappy interpretation of *if* is far-flung, the trouble is that the introduced gappiness ruins so much else. In order to get the right behavior for modals (and quantifiers and hedges) we have to reign in their quantificational force. For instance: instead of surveying points and checking whether they all make the prejacent true, we instead look to see if they all make the prejacent true-if-defined. This is a general pattern. Everywhere the standard set-up employs <u>true</u>, this Belnapian one has to employ <u>true-if-defined</u>.

An example: what we're up to in asserting stuff is somehow truth-directed. You try to assert only true things. (We don't want to pick sides in the various norms of assertion battles.) But assuming a story along the lines of Analysis 5, this will need some amending for indicatives. That's because we can assert conditionals with the hopes that their antecedents will turn out false:

(24) If you spill the beer, we will be all out.

So don't spill it! The fix here mirrors the fix for the modals: try to assert only true-if-defined things.

OK, where else does the standard set-up use <u>true</u>? Well, it plays a fairly prominent role in saying what entailment amounts to. In the classical set-up,

$p_1, \ldots, p_n \models_C q$ just in case all the worlds where all the p_i's are all true are worlds where q is true. So, in the Belnapian set-up, the thing to do is to reign in talk of truth to true-if-defined. This is a Strawson Entailment relation.[27]

Definition 6 (Strawson Entailment)**.** $p_1, \ldots, p_n \models_S q$ iff if $[\![p_1]\!]^{c,j}$ is defined and . . . and $[\![p_n]\!]^{c,j}$ is defined then if $[\![p_1]\!]^{c,j} = 1$ and . . . and $[\![p_n]\!]^{c,j} = 1$, then $[\![q]\!]^{c,j} = 1$ if defined.

If your ps and qs aren't gappy, this wrinkle adds nothing, which is good. But the definedness check blocks some otherwise bad entailments: for instance $\neg p \not\models_S$ *if* p, q.

The immediate problem though is that this means that conditionals entail (presuppose, actually) their antecedents, which ain't true (or even true-if-defined). The deeper problem is that the Belnapian story has coopted undefinedness — a tool with a well-marked out and understood use for marking and tracking presupposition failure — for its own purposes and left the cupboard bare for what we should use to replace it. #cricket-chirping.

There may be a dodge nearby. The dynamic machine behind Analysis 4 is expressive enough to capture the Belnapian proposal as a special case. Here is the easiest way to see that. First: require that an *if* (issued in a context) presupposes that the antecedent be compatible with that context. Second: evaluate the conditional for truth at limit-case states and collect up the results.

Definition 7 (Gap Retrieval)**.** Let S be any state and i any world in C_S.

i. $S[(\textit{if } p)(q)]$ is defined only if $C_{S[p]} \neq \emptyset$

ii. $\mathrm{Bel}(p) = \{i \in C\colon\ \{i\} + p = \{i\}\}$

The Belnap-profile (relative to C) of p is the result of (i) ignoring probabilistic information, (ii) updating the limit-case singleton states $\{i\}$ with p, and (iii) collecting up the worlds that survive. Most of the time this does nothing interesting since $\mathrm{Bel}(p) = C \cap P$ for most p. But not if p is gappy: just because $i \notin \mathrm{Bel}((\textit{if } p)(q))$ it doesn't follow that $i \in \mathrm{Bel}(\neg(\textit{if } p)(q))$ since $S[(\textit{if } p)(q)]$ can be undefined.[28]

So the dodge:

27 See von Fintel 1999.

28 Starr 2010 argues for a role for retrieved gaps in this spirit.

Analysis 6 (Belnap-on-demand)**.** Let $S(i) = \langle \{i\}, \mathrm{Prob}_S \rangle$. Then:

$$S \models \mathit{probably}\ (\mathit{if}\ \ p)(q) \ \text{ iff } \frac{\mathrm{Prob}(\mathrm{Bel}\,(\mathit{if}\ \ p)(q))}{\mathrm{Prob}(\{i\colon\ S(i)[(\mathit{if}\ \ p)(q)] \text{ is defined }\})}$$

is high enough.

Of course for this to work we have to wide scope the hedges.[29]

There is no technical issue here, but there is a not-small WTF-issue.[30] Why would most of the language (modals included!) want and care about the standard (dynamic) meaning of *if* but hedges want this very different function to operate on? The Belnap-profiles of modals are close to useless, so what motivates splitting the connection between modals and conditionals? Good questions. We'll move on.

7 High Types

That brings us to the last strategy we'll look at on our ascent. So far all hands are agreed that *if*s and operators — modals and hedges in our case — intermingle with ease and the result of said intermingling is that the operators get restricted by the *if*s. Taking *if* to be a restrictor (Analysis 3) accomplished this by taking *if* not to contribute a meaning of its own at all (bare conditionals being (covert) necessity modals). Taking *if* to be a dynamic conditional (Analysis 4) accomplished this by evaluating the consequent of the conditional in a subordinate context (while in effect modalizing the consequent). The first option takes the *if*-construction to be a modal, the second takes it to be a conditional.

There's a hybrid possibility here. Maybe the reason why *if*s and operators seem made for each other is that what *if* constructions mean is neither a modal nor a conditional but a function from (the semantic value of) a modal or hedge to a standard conditional meaning. As before, there are two (not quite equivalent) implementations of the idea.

First implementation: take the restrictor analysis and re-Schönfinkel/re-Curry the function. That is:

29 Had we opted for a less explicit version of the dynamic conditional, then gap retrieval amounts to evaluating p at diagonal, limit-case contexts and collecting the results: $\mathrm{Bel}(p) = \{i \in C : [\![p]\!]^{\{i\},i} = 1\}$. And then we insist that hedges care about this instead of the standard meaning.

30 Why That Function?

Analysis 7 (High-Type, Take #1)**.** Let $\mathbb{O}$ be a variable ranging over the semantic values of (binary) modals and hedges. Then:

$$[\![(if\ p)(q)]\!]^{c,i} = \lambda\mathbb{O}.\mathbb{O}([\![p]\!])([\![q]\!])$$

Bare conditionals, as you'd expect, take a covert $[\![\Box]\!]$.

(15) a. Alex: If he didn't tell Harry, he told Tom.
$\lambda\mathbb{O}.\mathbb{O}(P)(Q)([\![\Box]\!])$

b. Billy: Probably (so)./That's likely.
$\lambda\mathbb{O}.\mathbb{O}(P)(Q)([\![\mathit{probably}]\!])$

Since there is a logical constituent contributed by *if* it's not surprising that Billy's anaphors can pick it out. And — as is made plain by the λ-conversions — hedging that constituent can straightforwardly express the restricted meaning.

Second implementation: take the dynamic analysis and similarly lift its type so that it too wants an operator.

Analysis 8 (High-Type, Take #2)**.** Let $\mathbb{O}$ be a variable ranging over the semantic values of (unary) modals and hedges. Then:

$$S[(if\ p)(q)] = \lambda\mathbb{O}.\begin{cases} S & \text{if } S[p] \text{ is a fixed-point of } (\mathbb{O}[q]) \\ \mathbf{0} & \text{otherwise} \end{cases}$$

Bare conditionals, as you'd expect, take a covert $[\top]$ or $[\Box]$ (doesn't matter which, really).

(15) a. Alex: If he didn't tell Harry, he told Tom.
$\lambda\mathbb{O}.\left\{\begin{matrix} S & \text{if } S[p] \text{ is a fixed-point of } (\mathbb{O}[q]) \\ \mathbf{0} & \text{otherwise} \end{matrix}\right\}([\Box])$

b. Billy: Probably (so)./That's likely.
$\lambda\mathbb{O}.\left\{\begin{matrix} S & \text{if } S[p] \text{ is a fixed-point of } (\mathbb{O}[q]) \\ \mathbf{0} & \text{otherwise} \end{matrix}\right\}([\mathit{probably}])$

Just as before: Billy can pick out what Alex's *if* contributes — a high-type object — and hedge it. Since $S[p]$ is a fixed-point of $([\mathit{probably}][q])$ iff $S[p] \models \mathit{probably}\, q$ Billy thereby expresses the restricted meaning.

The two implantations are close relatives but not quite equivalent. Analysis 7 doesn't easily take right-nested *if*s but Analysis 8 isn't so sensitive.

There's something attractive to the high-type idea. (Beyond that it gets the restricted meanings to fall out.) In their original versions, the restrictor analysis and the dynamic analysis have a certain complementary relationship: one has modals doing conditionals' work (Analysis 3), the other going the other way (Analysis 4). The high-type view instead returns the work: modals and conditionals end up doing the work of modals and conditionals (respectively).

8 Parting Shots

There are different ways of making progress. One way lays out what everyone knows is a problem and provides a clean solution to it. That isn't the kind of progress we aim to make here because that isn't the kind of problem we currently face. Instead, we wanted to advertise a problem and make clear that it is worse than you might have expected or hoped.

References

Adams, Ernest W. 1975. The Logic of Conditionals. D. Reidel.

Belnap, Nuel D. 1970. Conditional assertion and restricted quantification. Noûs 4(1). 1–12.

Bennett, Jonathon. 2003. A philosophical guide to conditionals. Oxford University Press.

van Benthem, Johan. 1986. Essays in Logical Semantics. D. Reidel.

Bradley, Richard. 2000. A preservation condition for conditionals. Analysis 60(3). 219–222. http://dx.doi.org/10.1093/analys/60.3.219.

Edgington, Dorothy. 1995. On conditionals. Mind 104(414). 235–329. http://dx.doi.org/10.1093/mind/104.414.235.

Edgington, Dorothy. 2009. Conditionals. In Ed Zalta (ed.), Stanford encycolpedia of philosophy, vol. Spring 2009 Edition. URL http://plato.stanford.edu/archives/win2003/entries/conditionals/.

von Fintel, Kai. 1997. Bare plurals, bare conditionals, and *only*. Journal of Semantics .

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

von Fintel, Kai. 2003. Modals and conditionals revisited. URL http://web.mit.edu/fintel/fintel-2003-umass-epistemics.pdf. Colloquium, UMass Amherst.

Gärdenfors, Peter. 1982. Imaging and conditionalization. The Journal of Philosophy 79. 747–760.

Gibbard, Allan. 1981. Two recent theories of conditionals. In William Harper, Robert Stalnaker & Glenn Pearce (eds.), Ifs, 211–248. D. Reidel.

Gillies, Anthony S. 2004. Epistemic conditionals and conditional epistemics. Noûs 38. 585–616. http://dx.doi.org/10.1111/j.0029-4624.2004.00485.x.

Gillies, Anthony S. 2009. On truth-conditions for *If* (but not quite only *If*). The Philosophical Review 118(3). 325–348. http://dx.doi.org/10.1215/00318108-2009-002.

Gillies, Anthony S. 2010. Iffiness. Semantics and Pragmatics 3(4). 1–42. http://dx.doi.org/10.3765/sp.3.4.

Huitink, Janneke. 2009. Domain restriction by conditional connectives. URL http://semanticsarchive.net/Archive/zg2MDM4M/Huitink-domainrestriction.pdf.

Jackson, Frank. 2006. Indicative conditionals, revisited.

Kratzer, Angelika. 1981. The notional category of modality. In H. J. Eikmeyer & H. Rieser (eds.), Words, worlds, and contexts. new approaches in word semantics, 38–74. Berlin: de Gruyter.

Kratzer, Angelika. 1986. Conditionals. Chicago Linguistics Society 22(2). 1–15.

Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), Semantics: An international handbook of contemporary research, 639–650. Berlin: de Gruyter.

Kratzer, Angelika. 2012. Modals and conditionals. Oxford University Press.

Lewis, David. 1975. Adverbs of quantification. In E. Keenan (ed.), Formal semantics of natural language. Cambridge University Press.

Lewis, David. 1976. Probabilities of conditionals and conditional probabilities. The Philosophical Review 85. 297–315.

Moss, Sarah. 2015. On the semantics and pragmatics of epistemic vocabulary. Semantics & Pragmatics 8(5). 1–81. http://dx.doi.org/10.3765/sp.

Rothschild, Daniel. 2012. Do indicative conditionals express propositions? Noûs (to appear).

Spohn, Wolfgang. 1988. Ordinal Conditional Functions: A Dynamic Theory of Epistemic States. In William Harper & Brian Skyrms (eds.), Causation in Decision, Belief Change, and Statistics, vol. 2, 105–134. Reidel.

Stalnaker, Robert. 1984. Inquiry. MIT Press.

Starr, William. 2010. Conditionals and questions. Ph.D. thesis, Rutgers University.

Veltman, Frank. 1985. Logics for Conditionals. Ph.D. thesis, University of Amsterdam.

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

Yalcin, Seth. 2012. Bayesian expressivism. 133rd Session of the Aristotelian Society.