

Embedded Attitudes*

Kyle Blumberg and Ben Holguín

September 2018

Abstract

This paper presents a puzzle involving embedded attitude reports. We resolve the puzzle by arguing that attitude verbs take *restricted* readings: in some environments the denotation of attitude verbs can be restricted by a given proposition. For example, when these verbs are embedded in the consequent of a conditional, they can be restricted by the proposition expressed by the conditional’s antecedent. We formulate and motivate two conditions on the availability of verb restrictions: (i) a constraint that ties the content of restrictions to the “dynamic effects” of sentential connectives and (ii) a constraint that limits the availability of restriction effects to present tense verbs with first-person subjects. However, we also present some cases that make trouble for these conditions, and outline some possible ways of modifying the view to account for the recalcitrant data. We conclude with a brief discussion of some of the connections between our semantics for attitude verbs and issues concerning epistemic modals and theories of knowledge.

1 A puzzle

Let us start with a case:

Bill’s Holiday: Chris, Andrew, and I are discussing the details of Bill’s holiday this summer. We all know that Bill usually says goodbye before embarking on a trip. Chris says ‘I think that Bill is going to Costa Rica next week’; then Andrew says ‘Actually, I heard that Bill left for Cuba today’. I think for a moment, then utter (1):

*Forthcoming at *Journal of Semantics*. Earlier versions of this paper were presented at the International Conference on Truth, Logic and Philosophy at Peking University, at a semantics seminar led by Philippe Schlenker at NYU, and at departmental seminars at the University of Stellenbosch and the University of Cape Town. We would like to thank all of the participants at those presentations for their feedback. We would also like to thank Chris Barker, Jeremy Goodman, Sarah Moss, Philippe Schlenker, and Seth Yalcin for helpful discussion of various points. Harvey Lederman, Matt Mandelkern, Jake Nebel, Jim Pryor, Stephen Schiffer, and Trevor Teitel read earlier versions of the paper, and their comments improved the final product significantly. Finally, we are particularly grateful to Cian Dorr, Simon Goldstein, Daniel Rothschild, and two anonymous reviewers for their extensive feedback over the course of the project.

(1) If Bill is on a plane to Cuba, then I am surprised that he left without saying goodbye.

(1) is acceptable in the context of *Bill's Holiday*. But suppose that unbeknownst to any of us (2) is true:

(2) Bill is on a plane to Cuba.

Taking (1) at face value, (3) follows by modus ponens:

(3) I am surprised that Bill left without saying goodbye.

However, it is difficult to see how (3) could be true in our scenario. For one thing, I appear to lack any of the standard phenomenology associated with surprise. For another, it is plausible that a report 'S is surprised that P' is true only if 'S knows that P' is.¹ But given the circumstances it is not at all clear how I could know whether Bill has left without saying goodbye. So we have a case in which modus ponens seems to lead from true premises to a false conclusion.² Other logical rules, e.g. modus tollens, also appear to be threatened: it is not permissible to infer that (2) is false from (1) along with the unacceptability of (3).

Importantly, even if—despite all appearances—(3) is true in our scenario, conditionals similar to (1) raise a further puzzle. Suppose that Jane joins our discussion in *Bill's Holiday*. We all know she is friends with Bill, but we also know that Bill is more likely to tell one of us about a vacation to Cuba than he is to tell Jane (suppose he is better friends with us than with her). After hearing about our discussion, Jane says 'Oh, I know where Bill is right now. I talked to him on the phone this morning about his plans'. In this context, (4) is acceptable:

(4) If Bill is on a plane to Cuba, then I am surprised that Jane knows this but I don't.

So too is (5):

(5) If Bill is on a plane to Cuba, then I am surprised that he departed without my knowing.

However, if surprise reports are knowledge entailing in the sense mentioned above, then the consequent of each of (4) and (5) entails something which seems to be of the form 'I know that: P and I don't know that P', which itself entails the straightforwardly contradictory 'I know that P and I don't know that P' (taking for granted that knowledge is factive and distributes over conjunction). It is generally assumed that indicative conditionals with

¹We will often use normal quotes where, strictly speaking, corner quotes should be used. No confusion should arise.

²Some might be inclined to think that (1) really expresses a subjunctive, e.g. 'If Bill *were* on a plane to Cuba, then I *would be* surprised that he left without saying goodbye'. However, the inference from 'If Bill were on a plane to Cuba, then I would be surprised that he left without saying goodbye' and 'Bill is on a plane to Cuba' to 'I am surprised that Bill left without saying goodbye' is generally considered to be valid (Bennett, 2003). So the original problem remains.

consequents that are known to be false are unassertable—especially when the antecedents of those conditionals are *not* known to be false—so (4) and (5) are predicted to be unacceptable, contrary to fact.

The conditionals above featured the verb ‘surprise’, but as far as we can tell the same puzzles arise with virtually *all* attitude verbs. For instance, it is quite straightforward to construct similar cases with emotives such as ‘hope’ and ‘regret’:

Tennis: I have been teaching Chris tennis for the last six months and I know that he loves playing. One day my friend Jane reports that she saw someone in the distance injure themselves badly on a tennis court. Hearing this information—though with no particular reason to think it was Chris that Jane saw—I utter (6)/(7):

- (6) If Chris injured himself horribly on the tennis court, then I regret that I ever taught him how to play.
- (7) If Chris injured himself horribly on the tennis court, then I hope that he won’t blame me for his injuries.

Both (6) and (7) are acceptable in the context of *Tennis*, but suppose that unbeknownst to Jane or me (8) is true:

- (8) Chris injured himself horribly on the tennis court.

(9) and (10) then seem to follow by modus ponens:

- (9) I regret that I ever taught Chris how to play.
- (10) I hope that Chris won’t blame me for his injuries.

However, once again it is hard to see how either report could be true, since I seem to have no evidence that anything untoward has happened to Chris. Also, if it is shared knowledge that Chris always plays tennis in a secluded area with Andrew, who tends to panic in emergencies, we can say things like (11):

- (11) If Chris injured himself on the tennis court, then I regret that Andrew is the one who knows it.

As before, the consequent of (11) appears to be contradictory on standard semantics for attitude verbs, yet the conditional remains perfectly acceptable.

A similar phenomenon arises with doxastic verbs like ‘think’ and ‘suspect’. Consider (12) in the context of *Bill’s Holiday*, and (13) in the context *Tennis*:

- (12) If Bill is on a plane to Cuba, then I suspect that he is traveling first-class.
- (13) If Chris injured himself on the tennis court, then I think that Andrew is panicking right now.

Each conditional has a true reading in their respective contexts. However, if Bill really is on a plane to Cuba then it doesn't seem to follow that I suspect he is traveling first class. Similarly, if Chris really did injure himself, it doesn't seem to follow that I believe Andrew is panicking.

The puzzle also arises in constructions other than conditionals. Consider the following scenario, as well as the disjunctions that follow it:

Party: I have been looking forward to Ted's party for a while, and expect a lot of people to be there. However, upon arriving I only see a handful of people milling around the drinks table. Then I say:

- (14) Either a lot of people are on the deck outside, or I'm surprised that there are so few people here.
- (15) Either a lot of people are on the deck outside, or I regret that I didn't bring more friends.
- (16) Either a lot of people are on the deck outside, or I think I should have stayed at home.

(14)-(16) are acceptable in context, and raise similar issues to the conditionals above. If there *aren't* a lot of people on the deck outside then, e.g. 'I'm surprised that there are so few people here' follows from (14) by disjunctive syllogism. But intuitively it is a live possibility for me that there are a lot of people outside, in which case I do *not* know that there are few people at the party. So we have a case in which disjunctive syllogism seems to lead from true premises to a false conclusion.

Where Φ is an attitude verb, we will call sentences of the form 'If P, then S Φ s that Q' *attitude conditionals*. Each of the conditionals we have considered so far is an attitude conditional. We will call the subclass of attitude conditionals that appear to entail sentences of the form 'If P, then S Φ s that (P and S doesn't know that P)' *Fitch conditionals*.³ (4), (5), and (11) are all Fitch conditionals. We will also call sentences of the form 'Either P, or S Φ s that Q' *attitude disjunctions*. Finally, we will call the class comprised of attitude conditionals and attitude disjunctions *attitude constructions*.

The puzzles raised above show that an explanation of what is happening with attitude constructions is in order. This paper's aim is consider some of the difficulties involved in

³The name is suggested by the 'paradox of knowability' that was discovered in response to the work of Frederic Fitch (Brogaard & Salerno, 2013). The paradox features claims of the form 'It is possible that S knows that (P and S does not know that P)'.

offering a systematic explanation, but also to try to develop the approach we find most promising. The view we defend is that the epistemic and doxastic bases that feature in the denotations of attitude verbs are sometimes intersected with, or *restricted* by, a given proposition. For instance, on the natural reading of (1), ‘surprise’ as it appears in the consequent is restricted by the proposition expressed by the antecedent of the conditional, namely that Bill is on a plane to Cuba. We show how this resolves the puzzles introduced above. However, this response also raises further issues that are less easily addressed, the most important of which is *how* these restrictions get determined in each case. We show that there are challenges to providing a satisfying answer to this question. Indeed, our inquiry will ultimately raise more questions than it answers. Nevertheless, we hope to open up an interesting area for future research.⁴

The paper is structured as follows. In §2 we argue against the view that scopal ambiguities explain the behavior of attitude constructions. In §3 we put forward our preferred approach to the puzzles of §1: the view that attitude verbs sometimes exhibit restricted readings. We explain how appealing to restriction vindicates our intuitions about attitude constructions as well as the validity of, e.g., modus ponens. §§4-5 formulate and motivate two conditions on the availability of restrictions. The first condition connects the availability of non-trivial restriction to the dynamic properties of the logical connectives. The second puts constraints on the subject term of the report, as well as the tense of the attitude verb. Then in §6 we consider some data that seem to pose a problem for the conditions presented in the previous sections. §7 draws some morals for future semantic theorizing, while §8 concludes.

⁴As far as we are aware, Drucker (2017) and Jerzak (forthcoming) are the only other theorists who have explicitly discussed attitude conditionals and some of the problems that they raise. Drucker appeals to attitude conditionals such as (6) and (7) in the course of arguing for a certain type of radical externalism about non-doxastic attitudes. He accepts modus ponens and takes these conditionals at face value, concluding that the facts about our attitudes of, e.g., regret can be deeply external to us. Drucker does not discuss attitude disjunctions, Fitch conditionals, or attitude conditionals featuring verbs other than emotives. Indeed, it is important for Drucker’s arguments that there not be comparable conditionals with doxastics, and thus data such as (12) and (13) pose a problem for the view he defends (p.8). Moreover, Drucker’s general approach cannot handle Fitch conditionals, and it doesn’t carry over to the data presented in §6 concerning, e.g. ‘know’.

As for Jerzak, he is primarily concerned with attitude conditionals that feature ‘want’ (see §6 for a discussion of some of Jerzak’s data). He provides a semantics for ‘want’ that is “information sensitive”: want reports depend for their interpretation on a shiftable information state parameter. Jerzak does not consider attitude disjunctions, Fitch conditionals, or attitude constructions involving factive verbs. He does explicitly discuss attitude conditionals featuring ‘believe’, but argues that the attitude verb should be interpreted “wide-scope” with respect to the conditional Jerzak (forthcoming, 10). See §2 for further discussion of this response and some problems with it. We do not have the space to consider Jerzak’s semantics in detail here. But it is worth noting that we are sympathetic to some of the central aspects of Jerzak’s approach, namely that that ‘want’ is sometimes subject to restriction effects. However, it is not clear to us how the account could be carried over to other verbs, e.g. factives and doxastics. So, in short: although Drucker and Jerzak discuss attitude constructions and make a number of insightful observations about their properties, our examination is more detailed, and our perspective on their significance is different.

2 Wide-scoping

We begin with a relatively conservative response to the puzzles raised by attitude constructions. This response tries to explain the data in terms of scope. The idea, roughly, is that a sentence whose surface form is ‘If P, then $S \Phi$ s that Q’ or ‘Either P, or $S \Phi$ s that Q’ is ambiguous between two readings, whose logical forms we can represent as follows:

(17) a. If P, then ($S \Phi$ s that Q).

b. $S \Phi$ s that (if P, then Q).

(18) a. Either P, or ($S \Phi$ s that Q).

b. $S \Phi$ s that (either P, or Q).

(17a)-(18a) represent the “narrow-scope” reading of the conditional/disjunction, on which the attitude verb Φ takes scope only over Q. (17b)-(18b) represent the “wide-scope” reading, on which the verb takes scope over the entire conditional ‘If P, then Q’ or disjunction ‘Either P, or Q’.

One might think that resolving this ambiguity will help with our puzzles. For instance, it might be maintained that (12) (‘If Bill is on a plane to Cuba, then I suspect that he is traveling first-class’) should be read wide-scope, and that its logical form can be more perspicuously expressed by (19):

(19) I suspect that: if Bill is on a plane to Cuba, then he is traveling first-class.

‘I suspect that Bill is traveling first-class’ does not follow from the conjunction of (12) and (2) (‘Bill is on a plane to Cuba’), so this move does block the problematic inference.

However, there are a number of problems with this response. For one, some attitude constructions are simply not amenable to it. Consider (20):

(20) If Bill is on a plane to Cuba, then the person who I think he’s traveling with is Mary.

Relative clauses are usually taken to be “scope islands” for movement, which means that it is unclear how ‘think’ could take wider scope than the conditional (May, 1985). But (20) raises all the same worries as the other examples under consideration.

For another, even when a wide-scope interpretation is available, it often fails to provide the right reading of the relevant attitude construction. Consider (21):

(21) If Bill is on a plane to Cuba, then he is traveling first class and I suspect that he is drinking champagne.

(21) is assertable only if I am sure that Bill is traveling first class, given that he is on a plane to Cuba; but ‘I suspect that: if Bill is on a plane to Cuba, then he is traveling first-class and drinking champagne’ is assertable even when I have only a suspicion that Bill is traveling first class, given that he is on a plane to Cuba. So, these two sentences don’t seem to have the same meaning.⁵

A related problem is brought out by attitude constructions like (1) (‘If Bill is on a plane to Cuba, then I am surprised that he left without saying goodbye’): if this is read wide-scope, then its logical form can be represented by (22):

(22) I am surprised that: if Bill is on a plane to Cuba, then he left without saying goodbye.

The trouble is that (22) is not only false, but obviously so. It is not in the least bit surprising that if Bill is on a plane to Cuba, then he left without saying goodbye. At no point since Bill missed his chance to say goodbye has my evidence favored the proposition that if Bill is on a plane to Cuba, then he *did* say goodbye before leaving. This proposition is thus not something I could find surprising.

Finally, as Drucker (2017, 13) notes, attitude conditionals in which the argument to the verb is a propositional anaphor raise a particularly sharp challenge to the wide-scope response:

(23) If Bill is on a plane to Cuba, then that surprises me.

Granted the plausible assumption that the semantic contribution of ‘that’ in (23) is the proposition expressed by the antecedent, the wide-scope reading of (23) is equivalent to (24):

(24) I am surprised that: if Bill is on a plane to Cuba, then Bill is on a plane to Cuba.

This is clearly not the intended reading of (23). Although many things may surprise me, the tautology that if Bill is on a plane to Cuba, then he is on a plane to Cuba is not one of them.

The same points apply *mutatis mutandis* to the wide-scope interpretations of the Fitch conditionals and attitude disjunctions from §1. We conclude that a different style of explanation is called for.⁶

⁵Thanks to an anonymous reviewer for suggesting examples like (20) and (21). The reviewer also notes that examples analogous to (21) are used by Yalcin (2012) in arguments concerning the scope of embedded probability operators.

⁶Drucker (2017, 13-14) considers a different response on which attitude conditionals like ‘If P, then S Φs that Q’ are systematically reinterpreted along the lines of: ‘If S *finds out* that P, then S *will* Φ that Q’. We find Drucker’s objections to this reply compelling. Among other things, he points out that “deathbed sentences” like (25) are inhospitable to the reinterpretation strategy:

(25) If this is the last thought I have before dying, then I’m glad it’s such a philosophical one.

We concur, and add that attitude conditionals like the following prove similarly difficult:

3 Restriction

We suspect that the puzzling behavior of the attitude constructions in §1 is due to the semantics of their *attitude verbs*, rather than, e.g. the syntactic properties of attitude constructions. Indeed, the existence of Fitch conditionals is itself strong *prima facie* reason to suspect that the puzzle arises there. Fitch conditionals do not have the semantic phenomenology of conditionals with contradictory consequents. But if the verbs that appear in their consequents were to be interpreted uniformly, it would be difficult to explain why this is so. Thus, we favor the strategy of positing a multiplicity of readings for attitude verbs. In particular, we argue that attitude verbs can take their “normal” readings as well as a range of (soon to be specified) “special” readings, readings whose existence explains the puzzling behavior of attitude constructions.

By positing a multiplicity of readings we get a straightforward defense of the validity of, e.g., *modus ponens*. The idea is just that the “counterexamples” to the inference rules that we considered in §1 equivocate. For instance, (1) (‘If Bill is on a plane to Cuba, then I am surprised that he has left without saying goodbye’) is true on its most natural interpretation because ‘surprise’ as it appears in the consequent takes a special reading. In contrast, (3) (‘I am surprised that Bill has left without saying goodbye’) is false on its most natural interpretation because ‘surprise’ as it appears here, i.e. when the report is *unembedded*, tends to take a normal reading. But holding readings fixed—that is, keeping things uniformly special or uniformly normal—the inference is valid: if (1) and (2) are both true, (3) must be true as well. It’s just that we tend not to hold readings fixed when arguments like these are considered. Similarly for arguments involving (14) (‘Either a lot of people are on the deck outside, or I’m surprised that there are so few people here’) using disjunctive syllogism.

As for Fitch conditionals such as (5) (‘If Bill is on a plane to Cuba, then I am surprised that he has departed without my knowing’), the idea is that their natural interpretations are ones on which the wide-scope verbs, e.g. ‘surprise’, take special readings while the narrow-scope verbs, e.g. ‘know’, take normal readings. On these interpretations, a Fitch conditional’s consequent is non-contradictory, vindicating the intuition that Fitch conditionals are coherent.

We will develop these ideas more explicitly over the course of this section. But before we do we want to be explicit about our aims at this point in the paper. Our response to the puzzles of §1 outlined directly above brings with it two important questions: (a) what, semantically, is the difference between normal and special readings, i.e. what do the entries for attitude verbs need to look like in order for them to take both kinds of readings?; and (b) when are the special readings available, and what determines their content? In this

(26) If Bill’s on a plane to Cuba but has made sure that I never find out about it, then I am surprised that he is so secretive.

section, we only try to answer question (a). §§4-6 are devoted to trying to answer question (b), which is more challenging.

3.1 Restricted belief

Due to its simplicity we begin with an account of the special readings of the verb ‘believe’. We turn to more complicated verbs like ‘regret’ once the basic machinery is in place.

Let us suppose that for any given belief state, there is a unique, maximal set of possible worlds consistent with it. More concretely, let us say that for any subject S and world w , $\text{Dox}_{w,S}$ is the maximal set of worlds compatible with what S believes in w —i.e. S ’s *belief set* in w (Hintikka, 1962). The orthodox semantics for ‘believe’ can be presented as follows (‘ p ’ denotes the proposition expressed by P):

(27) **Standard semantics for ‘believe’**

‘ S believes that P ’ is true at w iff: $\text{Dox}_{w,S} \subseteq p$.

Now we introduce a mechanism that, when applied to a set of possibilities, will allow us to produce a (possibly strict) subset of those possibilities. We will call this mechanism *restriction*. A restriction \dagger is a set of worlds, i.e. a proposition. A restriction can be used to winnow down a set of possibilities via set intersection. We allow restrictions to have a semantic effect by enriching our points of evaluation with a propositional parameter:

(28) **Restricted semantics for ‘believe’**

‘ S believes that P ’ is true at w under restriction \dagger iff: $(\text{Dox}_{w,S} \cap \dagger) \subseteq p$.

To illustrate, suppose that John’s belief state at w_1 can be represented by three worlds: $\text{Dox}_{w_1,\text{John}} = \{w_1, w_2, w_3\}$, and that the proposition $A = \{w_1, w_4\}$. Suppose also that Bill is in Costa Rica in w_1 , and the United States in w_2 and w_3 . Given (28), John believes that Bill is in Costa Rica’ is true at w_1 under A iff $(\text{Dox}_{w_1,\text{John}} \cap A) \subseteq \{w \mid \text{Bill is in Costa Rica at } w\}$. Since $(\text{Dox}_{w_1,\text{John}} \cap A) = \{w_1\}$ and $\{w_1\} \subseteq \{w \mid \text{Bill is in Costa Rica at } w\}$, it follows that $(\text{Dox}_{w_1,\text{John}} \cap A) \subseteq \{w \mid \text{Bill is in Costa Rica at } w\}$. So, the report is true at w_1 under the restriction A .⁷

Note that the standard semantics for ‘believe’ may be straightforwardly recovered from the restriction semantics. This is because whenever \dagger is the *trivial restriction* \top —the set of all worlds—it follows that for any subject S , and world w : $\text{Dox}_{w,S} \cap \top = \text{Dox}_{w,S}$. Thus, from the perspective of the restricted account, the reading of ‘believe’ posited by the standard account is just that in which the restriction is the tautology.

⁷For simplicity, we will assume that restrictions are not derived from more complex objects, e.g. as the intersection of a “modal base” (Kratzer, 1986), or the output of a question denotation given a relevant world. However, see fn.33 for further discussion of the latter option.

Before explaining how the restriction semantics can account for §1’s doxastic attitude constructions, it will be helpful to introduce a shorthand for referring to the various readings induced by restriction. We will enrich our metalanguage as follows. Where \dagger is a restriction, we will call $\text{Dox}_{w,S} \cap \dagger$ the set that is determined by S ’s *beliefs* $_{\dagger}$ (at w). Accordingly, if $(\text{Dox}_{w,S} \cap \dagger) \subseteq p$, then we will say that S *believes* $_{\dagger}$ that P (at w). Given the observation above, when the trivial restriction is in play we will speak interchangeably of *believing* and *believing* $_{\top}$.

To reiterate, our goal here is to give an explicit account of how special and normal readings *work*, not how special and normal readings *get determined*. So at this point we will just help ourselves to the relevant restrictions and show how our restricted semantics generates the desired readings. With that in mind, let I be the proposition that Chris injured himself on the tennis court. The natural readings of (13) (‘If Chris injured himself on the tennis court, then I think that Andrew is panicking right now’) and ‘I think that Andrew is panicking right now’ (i.e. the consequent of (13) as it occurs unembedded) can be represented as follows:

(29) I think $_{\top}$ that Andrew is panicking right now.

(30) If Chris injured himself on the tennis court, then I think $_I$ that Andrew is panicking right now.

Since think $_{\top}$ is just normal belief, and since I definitely do not believe in anything like the normal sense that Andrew is panicking, ‘I think that Andrew is panicking right now’ is false in *Tennis* when the constraint is \top .⁸ Thus, the unembedded report is false on its natural reading, just as we want.

As for (13), given that Chris did actually injure himself at $w_{\textcircled{a}}$ (the actual world), its semantic value hangs on whether ‘I think that Andrew is panicking right now’ is true at $w_{\textcircled{a}}$ under the restriction I . This is indeed the case, since every world in $\text{Dox}_{w_{\textcircled{a}},\text{Me}}$ in which Chris injured himself is one where Andrew panics. That is, the set of possibilities corresponding to the intersection of all that I think $_{\top}$ with the proposition that Chris injured himself on the tennis court is a subset of the set of possibilities in which Andrew is panicking right now. Hence, (13) is true on its most natural reading.

To sum up, we maintain that it is the non-uniformity of restriction in the natural interpretations of embedded and unembedded attitude reports that explains our seemingly inconsistent intuitions about attitude constructions featuring doxastics. The remainder of the section will offer a similar analysis of ‘know’ and ‘regret’. This will help to illustrate how the restriction account can be applied more generally.

⁸We assume harmlessly that ‘believe’ and ‘think’ are semantically equivalent (cf. Hawthorne *et al.* (2016)).

3.2 Factive verbs and (restricted) knowledge

‘Regret’ is a so-called “factive” attitude verb (Giannakidou, 2006). We will assume the following about factive verbs:

K-ENTAILING: Where Φ is a factive attitude verb: if ‘S Φ s that P’ is true, then ‘S knows that P’ is true.⁹

K-ENTAILING was already alluded to in the presentation of the puzzles in §1. For instance, we argued that part of the reason to think (3) (‘I am surprised that Bill left without saying goodbye’) is false is that, intuitively, I don’t *know* that Bill has left without saying goodbye, so I can’t be *surprised* that he left without saying goodbye either.

We endorse K-ENTAILING and intend to defend a restriction-based semantics for factives like ‘regret’. But this raises an issue: we know that the subjects of our various attitude constructions lack anything resembling normal knowledge of the propositions expressed by the complement clauses of the embedded reports. It follows that we need a semantics for ‘know’ that makes it amenable to the phenomenon of restriction too. Let us stipulate that for any subject S and world w , $Epi_{w,S}$ is the maximal set of worlds compatible with what S knows in w —i.e. S ’s *knowledge set* in w . Here are the standard and restricted entries for ‘know’:

(31) **Standard semantics for ‘know’**

‘S knows that P’ is true at w iff: $Epi_{w,S} \subseteq p$.

(32) **Restricted semantics for ‘know’**

‘S knows that P’ is true at w under restriction \dagger iff: $(Epi_{w,S} \cap \dagger) \subseteq p$.¹⁰

With these entries for ‘know’ in place, we can speak meaningfully of what a subject knows _{\dagger} , and thus make use of restricted knowledge in giving a semantics for verbs like ‘regret’. We should be clear, however, that the motivation for our restricted semantics for ‘know’ goes beyond the fact that it is required by the conjunction of K-ENTAILING and our semantics for emotive factive verbs like ‘regret’. As we will see in §6, there are a variety of attitude constructions involving ‘know’ that tell in favor of a semantics along the lines of

⁹See (Williamson, 2000, ch. 1) for arguments in support of this principle. Egré (2008) maintains that factives are only “belief entailing” so that, e.g. ‘S regrets that P’ only entails ‘S believes that P’. Those who follow him in this are welcome to weaken K-ENTAILING accordingly, for this will make no difference to our central arguments. Also, those who think that factivity is best captured as a presupposition rather than an entailment are welcome to substitute “undefined” for “false” in the relevant arguments. Again, this will make no difference to the central claims of the paper.

¹⁰It is worth noting that Schaffer & Szabo (2013) give an entry for ‘know’ on which it is essentially modeled as an adverb of quantification (like ‘usually’ or ‘might’). In particular, they allow the modal base that ‘know’ quantifies over to be restricted by the antecedent of a conditional. However, when discussing attitude conditionals featuring ‘know’, Schaffer & Szabo, (pp.528-30) maintain that ‘know’ takes wide-scope over the conditional. So, although their semantics allows for restriction effects, they do not make use of it in solving the problems posed by attitude constructions.

(32). But as we will also see in §6, this data is far more unruly than the data involving the other attitude verbs that we have considered so far. We thus postpone the task of offering direct motivation for (32) until later in the paper.

3.3 Restricted regret

For simplicity, we will take our baseline semantics for ‘regret’ to be Heim’s (1992) *comparative desirability* account.¹¹ The rough idea behind Heim’s theory is that regretting something is a matter of knowing it, but also wishing it weren’t so. A bit more formally:

(33) **Heim-style standard semantics for ‘regret’**

‘S regrets that P’ is true at w iff: (i) S knows p at w and (ii) for most worlds w' compatible with what S knows at w : S prefers (at w) the worlds most similar to w' in which $\neg p$, to w' .¹²

Thus, ‘John regrets that Bill failed the exam’ is true at w iff (i) John knows that Bill failed the exam, and (ii) for most worlds w' compatible with John’s knowledge: John prefers each world most similar to w' in which Bill did not fail the exam, to w' .¹³

To handle the attitude constructions of §1 that feature ‘regret’, we enrich Heim’s entry with restriction:

(34) **Restricted semantics for ‘regret’**

‘S regrets that P’ is true at w under \dagger iff: (i) S knows $_{\dagger}$ p at w and (ii) for most worlds w' compatible with what S knows $_{\dagger}$ at w : S prefers (at w) the worlds most similar to w' in which $\neg p$, to w' .

The natural readings of, e.g. (6) (‘If Chris injured himself horribly on the tennis court, then I regret that I ever taught him how to play’) and (9) (‘I regret that I ever taught him

¹¹We could just as easily have used an “ideal worlds” analysis (von Fintel, 1999), but a Heim-style account is more economical for our purposes. Strictly speaking, Heim does not explicitly provide an entry for ‘regret’, but it is easy to see what she intends given her treatment of ‘glad’.

¹²Two remarks about this entry are in order. First, by assuming Strong Centering, i.e. that if w is a p -world, then the closest p -world to w is w itself (Lewis, 1973), we are able to simplify condition (ii) from what Heim has, which officially is ‘... S prefers (at w) the worlds most similar to w' in which $\neg p$, to the worlds most similar to w' in which p ’ (although Heim does suggest that such a simplification is plausible). Second, the analog of (ii) in Heim’s entries for desire-based verbs quantifies *universally* over S ’s knowledge set, whereas our entry checks only to see whether a *majority* of the worlds in S ’s knowledge set meet the relevant condition. A Heim-style condition makes the problematic prediction that if it’s an epistemic possibility that your lottery ticket will turn out to be a winner, then absent bizarre preferences against winning lots of money, you can’t regret buying it. That said, neither (ii) nor the standard Heimian truth-conditions can capture so-called ‘insurance’ cases: ‘I regret that I didn’t buy house insurance’ can be true even when it is quite likely on my evidence that nothing happens to my house (Levinson, 2003). Levinson proposes a probabilistic semantics in response to these cases. As far as we can tell, this semantics is compatible with the mechanism of restriction that we posit, but we stick with something more straightforward.

¹³The reason only the worlds *closest* to w' in which Bill did not fail the exam are considered is that the report can be true even if John prefers some distant worlds in which Bill fails the exam to w' (e.g., worlds where all and only those who pass the exam are enslaved by cruel aliens).

how to play’) can be represented as follows (recall that I is the set of worlds in which Chris injured himself on the tennis court):

- (35) If Chris injured himself horribly on the tennis court, then I regret $_I$ that I ever taught him how to play.
- (36) I regret $_{\top}$ that I ever taught Chris how to play.

It is straightforward to show that on the restrictions assigned above, (6) is true but (9) is false. What is perhaps less obvious to see is how restriction helps us make sense of the Fitch conditional (11) (‘If Chris injured himself on the tennis court, then I regret that Andrew is the one who knows it’). Its most natural reading can be represented as follows:

- (37) If Chris injured himself on the tennis court, then I regret $_I$ that Andrew is the one who knows $_{\top}$ it.

That is, in order to handle (11) we assume that ‘regret’ and ‘know’ in the consequent are evaluated under distinct restrictions: I and \top , respectively. Given that Chris injured himself at $w_{\textcircled{a}}$, this means that (11) is true at $w_{\textcircled{a}}$ iff (i) at $w_{\textcircled{a}}$ I know $_I$ that Andrew is the one who knows $_{\top}$ that Chris is injured (i.e. for all $w \in \text{Epi}_{w_{\textcircled{a}}, \text{Me}}$ such that Chris injured himself in w , for all x except Andrew: $\text{Epi}_{w,x} \not\subseteq I$), and (ii) for each w' compatible with what I know $_I$ at $w_{\textcircled{a}}$: I prefer the worlds most similar to w' in which someone distinct from Andrew is the only person who knows $_{\top}$ that Chris injured himself, to w' . Condition (i) is satisfied given my knowledge $_{\top}$ that Andrew and Chris play tennis in a secluded area, i.e. my knowledge $_{\top}$ that if Chris was injured, only Andrew would know $_{\top}$ about it. Condition (ii) is also satisfied given my knowledge $_{\top}$ that Andrew tends to panic in emergencies. So, (11) is true at $w_{\textcircled{a}}$, as required.

Taking stock

We trust that the general form of the restriction strategy is now clear enough to see how it could be applied to other kinds of attitude constructions. Thus, one need not think that the puzzling behavior of attitude constructions is due to the invalidity of inference rules like modus ponens or disjunctive syllogism. Nor does one have to think that attitude constructions exhibit surprising syntactic properties. So long as one accepts that attitude verbs are subject to the mechanism of restriction, one will have the resources to account for the natural readings of the data in §1 in a manner that is logically and syntactically orthodox.

However, the picture is not complete, for we have not yet said when non-trivial restrictions take effect, or what determines their content when they do. The second half of the paper is devoted to answering these questions. For reasons that will emerge over the course of the discussion, we believe that the issues involved here are of considerable complexity. But we

also believe that one of the main contributions of the first half of this paper is that we are now in a position to understand these challenges more clearly.

In the following two sections, we formulate and motivate two conditions on non-trivial restriction. Both conditions draw on the *formal* properties of attitude constructions to constrain the availability and content of the restriction. The first condition connects the availability of non-trivial restriction to the dynamic properties of conditionals, disjunctions, and conjunctions (§4). The second puts constraints on the subject term of the report, as well as the tense of the attitude verb (§5). We argue that taken together, these restrictions provide a fairly satisfying picture of the data considered so far. (But to foreshadow what will come in §6, it is unclear how robust these generalizations are once the scope of the inquiry expands to attitude verbs such as ‘know’, ‘remember’, and ‘want’.)

4 Dynamic restriction

Recall that I is the set of worlds in which Chris injured himself on the tennis court, and let D be the set of worlds in which there are not a lot of people on the deck outside. We know that the (a) sentences take the (b) readings (in their respective contexts).

- (38) a. If Chris injured himself horribly on the tennis court, then I regret that I ever taught him how to play.
 b. If Chris injured himself horribly on the tennis court, then I regret _{I} that I ever taught him how to play.

- (39) a. I regret that I ever taught Chris how to play.
 b. I regret _{\top} that I ever taught Chris how to play.

- (40) If Chris injured himself on the tennis court, then I regret that Andrew is the one who knows it.

If Chris injured himself on the tennis court, then I regret _{I} that Andrew is the one who knows _{\top} it.

- (41_a). a. Either a lot of people are on the deck outside, or I regret that I didn’t bring more friends.
 b. Either a lot of people are on the deck outside, or I regret _{D} that I didn’t bring more friends.

What we want to know is *why*. That is, we want to be able to predict that when a regret report is embedded in a conditional or a disjunction, the relevant epistemic base can be

non-trivially restricted; but when it is unembedded, it seems that it cannot be. Moreover, we want the embedded report to take the right restriction, e.g. the proposition expressed by the antecedent of the conditional in (38a), and the negation of the first disjunct in (41a). We also want the restriction to be optional in order to handle the pattern of restrictions that we find in Fitch conditionals like (40). We will argue that a natural way of trying to achieve these results is by systematically tying restrictions to the “dynamic” properties of sentential connectives. We spell this out below.

It is commonplace to find theorists maintaining that sentential connectives have “dynamic effects”. The general thought can be illustrated through the phenomenon of presupposition projection. Consider the following sentences:

- (42) a. John stopped smoking last week.
 b. John started smoking last year and he stopped smoking last week.
 c. If John started smoking last year, then he’s stopped smoking this year.
 d. Either John never smoked, or he stopped smoking this year.

(42a) carries the presupposition that John used to smoke. However, none of (42b)-(42d) carry this presupposition. The standard explanation is that in each case, the presupposition is (dynamically) “filtered” by earlier material in the sentence. That is, the embedded, presupposition-carrying clause (‘John stopped smoking’) is evaluated against a background that is provided by previous parts of the relevant sentence. Since the presupposition that John used to smoke is satisfied in this enriched background, the sentence as a whole presupposes nothing. In (42b) this background is provided by the first conjunct, i.e. the proposition that John started smoking last year; in (42c) it is provided by the antecedent of the conditional (again the proposition that John started smoking last year); and in (42d) it is provided by the negation of the first disjunct, i.e. the proposition that it is not the case that John never smoked, i.e. that John used to smoke.¹⁴

Though we wish to remain neutral on how exactly these dynamic effects should be captured, and what exactly explains their presence, we believe that something in the vicinity can be used to place substantive constraints on the availability of non-trivial restriction. First, we assign clauses *dynamic propositions* as follows:

- (43) a. DYNAMIC PROPOSITIONS
 b. If π is a main clause, then π is assigned \top .

¹⁴There is an enormous literature on these sorts of dynamic effects and how they should be modeled. The idea that the interpretation of presupposition-carrying constituents in a sentence systematically depend on earlier material essentially goes back to Karttunen (1974) and Stalnaker (1974). This thought has been developed in both a semantic direction, e.g. (Heim, 1982), (Beaver, 2001), and a pragmatic one, e.g. (Schlenker, 2009). Klinedinst & Rothschild (2012) discuss such dynamic effects as they relate to a variety of constructions, e.g. modals and adverbs of quantification.

- c. (i) If a conjunction $\psi \wedge \chi$ is assigned R , then ψ is assigned R and χ is assigned $R \cap \llbracket \psi \rrbracket$
(ii) If a conditional $\psi \rightarrow \chi$ is assigned R , then ψ is assigned R and χ is assigned $R \cap \llbracket \psi \rrbracket$
(iii) If a disjunction $\psi \vee \chi$ is assigned R , then ψ is assigned R and χ is assigned $R \cap (\top - \llbracket \psi \rrbracket)$

Then we tie the possibility of non-trivial restriction to dynamic propositions:

(44) DYNAMIC RESTRICTION

Given a main clause π , and constituent ‘ $S \Phi$ s that P ’ of π : Φ is non-trivially restricted by \dagger only if $\dagger =$ the dynamic proposition assigned to ‘ $S \Phi$ s that P ’.

In short, what DYNAMIC RESTRICTION says is that if an attitude verb ever takes a non-trivial restriction \dagger (i.e. $\dagger \neq \top$), then \dagger must be equivalent to the background for interpretation that is systematically provided by previous parts of the sentence (as determined by the rule for the relevant connective).

Note that DYNAMIC RESTRICTION only allows *embedded* reports to take non-trivially restricted readings, since only embedded clauses are assigned non-trivial dynamic propositions (by (43b)). This explains why, e.g., (39a) (‘I regret that I ever taught Chris to play tennis’) is unacceptable, for its ‘regret’ can only be trivially restricted.

But things are different for embedded reports. For instance, by (ii) of (43c), ‘regret’ can be restricted by I in (38a) (‘If Chris injured himself on the tennis court, then I regret that I ever taught him how to play’). (38a) is true when ‘regret’ takes this restriction. A similar result is obtained in the case of (41a): if ‘regret’ is non-trivially restricted then, by (iii) of (43c), the restriction must be D . (41a) is true when the verb takes this restriction. So, when ‘regret’ takes a non-trivial constraint in (41a), the sentence is true. Thus, DYNAMIC RESTRICTION goes some way in explaining why (38a) and (41a) take their respective readings.

What about Fitch conditionals like (40) (‘If Chris injured himself on the tennis court, then I regret that Andrew is the one who knows it’)? Here we appeal to the fact that DYNAMIC RESTRICTION provides only a *necessary* condition on non-trivial restrictions, not a sufficient one. Even when attitude verbs are embedded, they need not take restricted readings—non-trivial restrictions are optional. This accounts for the pattern of restrictions that we find in Fitch conditionals, e.g. (40): ‘surprise’ is non-trivially restricted but ‘know’ is not.¹⁵

¹⁵We do not want to oversell the optionality of restriction, however. As far as we can tell, attitude constructions like (45) only have false readings:

(45) ?? If Bill is on a plane to Cuba, then although I don’t know that he left without saying goodbye, I am

Finally, it is worth observing that DYNAMIC RESTRICTION makes a striking prediction, namely that attitude constructions should exhibit order effects. For example, it predicts that ‘S Φ s that P’ must be the *second* disjunct in a disjunction if Φ is to take a non-trivial restricted reading. As far as we can tell this prediction is indeed borne out by the data:

- (46) a. Either a lot of people are on the deck outside, or I’m surprised by how few people there are here.
 b. Either I’m surprised by how few people there are here, or a lot of people are on the deck outside.
- (47) a. Either a lot of people are on the deck outside, or I regret that I didn’t bring more friends.
 b. ?? Either I regret that I didn’t bring more friends, or a lot of people are on the deck outside.

This provides yet further support for a condition on restriction along the lines of DYNAMIC RESTRICTION.^{16,17}

surprised that he left without saying goodbye.

But all that is needed for (45) to have a true reading is for ‘know’ to take the trivial restriction while ‘surprise’ takes a non-trivial restriction. So perhaps the following should be added as an independent structural constraint on the workings of restriction: whenever two or more attitude verbs take the same scope in a sentence, they must all take the same restriction (be it trivial or non-trivial). Otherwise the restrictions may vary. Attitude conditionals like (45) feature verbs that take the same scope; whereas Fitch conditionals feature verbs that vary in scope. Hence the difference in the availability of the non-uniform readings.

¹⁶One concern with DYNAMIC RESTRICTION is that the restrictions it makes available appear to be too strong in some cases. Consider (48):

- (48) If Chris injured himself on the tennis court, then I regret that I ever taught him how to play and I’m worried that he’ll never speak to me again.

Intuitively, the proposition that Chris injured himself (*I*) is a more plausible restriction for ‘worried’ than the proposition that Chris injured himself *and I regret_I that I taught him how to play*. That is, it seems that we want to have available a *weaker* restriction than the one that DYNAMIC RESTRICTION generates. We won’t provide a full solution to this problem here, but believe that it can be remedied by assigning clauses *sequences* of propositions, rather than just propositions. If this is done in the right way, the proposition that Chris injured himself will be a member of the sequence assigned to the ‘worry’ report in (48). One can then allow attitude verbs to be optionally restricted by an element in the sequence which they are assigned.

¹⁷One might wonder why we have not tied the possibility of restriction to *local contexts* (Karttunen, 1974; Stalnaker, 1974; Heim, 1982; Schlenker, 2009). The idea would be the following: if an attitude verb ever takes a non-trivial restriction \dagger , then \dagger must be equivalent to the local context of the relevant report. This proposal is similar in spirit to DYNAMIC RESTRICTION, but has the effect of strengthening the content of the restriction from the relevant dynamically supplied proposition to the intersection of that proposition and the global context in which the utterance occurs.

One problem with this proposal stems from the fact that interlocutors may mutually presuppose something false. In this case, the context set will contain false information and can lead to unwanted restricted readings. For instance, suppose that we all come to falsely believe that Bill is on a plane to Cuba, and then start discussing how odd it is that he left so abruptly. I say (49):

5 First-Present Restriction

So far we have discussed a condition that ties the possibility of non-trivial verb restriction to the environment in which the relevant report is embedded. Now we consider a condition that ties restrictions to the form of the attitude report itself.

5.1 Motivating the condition

Attentive readers may have noticed that the attitude reports under consideration have invariably been stated in the present tense using a first-person pronoun. This is no coincidence. The phenomenon of verb restriction is in general much easier to get when the attitude reports have these properties. Consider speeches like the following:¹⁸

(50) ?? I'm sleepy and Bill believes I'm sleepy.

(51) ?? It's raining and I was surprised it was going to rain today.¹⁹

If the only constraint on restriction was DYNAMIC RESTRICTION, then (50) would have a true reading so long as it is merely *compatible* with what Bill believes (in the normal sense) that I am sleepy. But this is clearly not right. If we know that for all Bill believes (in the normal sense) I am wide awake, then (50) seems only to have false readings. Likewise for (51): given only DYNAMIC RESTRICTION, the mere compatibility of my past knowledge state with the proposition that it would be raining today would be enough to make possible true, restricted readings of (51). But it seems only to have false readings.

These asymmetries in person and tense and can be seen in conditional and disjunctive attitude constructions as well.²⁰ Note that the past and future tense versions of (1) are

(49) I am surprised that Bill left without saying goodbye.

(49) only has a false reading (since the complement is false). However, the local context for the report *C* (i.e. the global context) contains false information that is not included in the epistemic states of the interlocutors, namely that Bill is on a plane to Cuba. This means that if 'surprise' in (49) is restricted by *C*, this will have a non-trivial effect. Specifically, it should be possible for 'surprise' in (49) to take a reading on which it is restricted by the proposition that Bill is on a plane to Cuba. But in that case, (49) should have a true reading, much in the same way that (1) ('If Bill is on a plane to Cuba, then I am surprised that he left without saying goodbye') has a true reading when 'surprise' in the consequent is restricted by the proposition expressed by the antecedent. In short, identifying restrictions with local contexts is problematic because the context set can potentially carry unwanted information. Thanks to an anonymous reviewer for bringing out this worry with the proposal.

¹⁸We also consider sentences of the form 'P and S Φs that Q' to be attitude constructions, namely *attitude conjunctions*.

¹⁹To be clear, (50)-(51) can have true readings. Our claim is only that these sentences cannot be read like the other attitude constructions discussed so far.

²⁰Thus, one should not take (50)-(51) to show that the phenomenon of verb restriction only arises in attitude conditionals and disjunctions. Notice as well that although (3) ('I am surprised that Bill left without saying goodbye') is intuitively false in the context of *Bill's Holiday*, the following disjunction of conjunctions is intuitively true:

(52) Either Bill is on a plane to Cuba and I'm surprised that he's left without saying goodbye, or he's at home

unacceptable:

- (53) a. ?? If Bill is on a plane to Cuba, then I *was* surprised that he left without saying goodbye.
b. ?? If Bill is on a plane to Cuba, then I *will be* surprised that he left without saying goodbye.

(53a)-(53b) seem false, or at the very least unassertable. And to the extent that we can get ourselves in a context where either seems true, such a context seems inevitably to be one in which ‘surprise’ takes its normal reading. For instance, if Bill is on a plane to Cuba and, e.g., (53b) is true, then it does seem to follow that I *will be* surprised (in the normal sense) that he left without saying goodbye.

The same is true (*mutatis mutandis*) of (54)-(55) in the context of *Party*:

- (54) ?? Either a lot of people are on the deck outside, or I was surprised that there would be so few people here.
(55) ?? Either a lot of people are on the deck outside, or I will regret not inviting more people.

In each case the disjunction either seems clearly false, or to be an expression of uncertainty about which normal attitudes I have had or will have.

With regard to the issue of person, suppose we add the following detail to *Bill’s Holiday*: we know that Amy, who is currently somewhere very far away, is just like us with respect to her expectations about and knowledge of Bill’s travel plans. In this context the following strike us as unassertable:

- (56) a. ?? If Bill is on a plane to Cuba, then Amy is surprised that he left without saying goodbye.
b. ?? Either Bill is on a plane to Cuba, or Amy is pleased that he is still at home.

Again, the inference in each case is that Amy has some special information about Bill’s whereabouts. But since it’s part of the setup of the case that she is as ignorant about his whereabouts as we are, (56a)-(56b) seem only to have false readings.

We take the above observations to provide a strong *prima facie* case for the following constraint on non-trivial restriction:

- (57) FIRST-PRESENT (FP) RESTRICTION
If Φ is non-trivially restricted by \dagger in ‘S Φ s that P’, then S is a first-person pronoun and Φ is in the present tense.

as expected and I’m happy that I’ll get to see him this evening.

But (52) would be false if ‘surprise’ and ‘happy’ could only take their normal readings.

It seems to us difficult to deny that restricted readings of embedded attitude verbs are easier to come by when they meet the requirements of FP RESTRICTION. However, unlike DYNAMIC RESTRICTION—which for the kinds of verbs considered so far in the paper seems to face few counterexamples—there are a variety of attitude constructions that make trouble for FP RESTRICTION. As such, we are unsure whether FP RESTRICTION should be taken to be a full-blooded constraint on the presence of non-trivial restriction, or merely a reliable generalization. The remainder of this section discusses two of these problematic cases, as well as some possible defensive moves available to those who would like to treat FP RESTRICTION as a genuine constraint on restriction.

5.2 Some problems for the condition

The first putative counterexample concerns “echoing” uses of attitude constructions.²¹ These arise when some subject, *S*, utters a present tense attitude construction in the first-person—say ‘If *P*, then I Φ that *Q*’—and then someone who overhears *S*’s speech echoes it in the third-person. So, for example: Amy asserts (1) (‘If Bill is on a plane to Cuba, then I am surprised that he left without saying goodbye’); we overhear it; then one of us asserts (56a) (‘If Bill is on a plane to Cuba, then Amy is surprised that he left without saying goodbye’). In these circumstances (56a) seems mostly fine, even when it is also salient to us that Amy’s evidence concerning Bill’s whereabouts is as impoverished as ours is. This seems to provide evidence that third-personal attitude constructions like (56a) sometimes take restricted readings, and thus that FP RESTRICTION constraint is false.

The second putative counterexample, due to an anonymous reviewer, concerns cases such as the following:

- (58) If Weyland Corporation’s earnings miss expectations, the stock price will fall and our financial advisor is of the opinion that we should buy.
- (59) If the bell rings, then the person who Bob suspects is at the door is Steve.

We share the reviewer’s judgment that these sentences can have true readings even when (i) it is not assumed that the subject of the embedded report knows whether the antecedents of the conditionals are true, and (ii) it seems like such knowledge should be required for the embedded belief report to be true. Moreover, the anti-wide-scoping arguments of §2 apply just as well to (58)-(59), so it is not plausible in either case that the doxastic is taking scope over the entire conditional. There is thus reason to think (58)-(59) exhibit the phenomenon of restriction, and thus that FP RESTRICTION is not true in full generality.

We are not entirely sure what to make of echoing uses of attitude constructions, or conditionals such as (58)-(59). They may well be genuine counterexamples to FP RESTRICTION.

²¹Thanks to Harvey Lederman and an anonymous reviewer for bringing such cases to our attention.

But there are also things that can be said in defense of the condition. Regarding echoing uses, it is worth observing that though one can hear a conditional such as (56a) sounding okay in these circumstances, just about any echoed speech can be made to sound okay, no matter its content. For example: suppose Amy is not part of our conversation, but we have an audio recording of her confidently uttering ‘I know Oswald didn’t shoot Kennedy’. If delivered in the right way, an utterance of ‘Amy knows Oswald didn’t shoot Kennedy’ would seem just fine. But it is clear that the function of the report is “quasi-quotational”, rather than straightforwardly assertoric. Perhaps something similar is driving the acceptability judgments involving echoing uses of third-personal attitude constructions.

As for the third-personal doxastic attitude conditionals such as (58) and (59), they seem to bear a similarity to the phenomenon of so-called “ultra-liberal” de re attitude reporting (Pryor 2004; Recanati 2012; Blumberg & Holguín 2018). This is when a subject who holds the purely general belief that all Fs are G can, in some contexts, be appropriately ascribed the particular belief that *o* is G, for some *o* that is in fact F. Importantly, this can happen even when it is known to the ascriber that, intuitively speaking, the subject of the report does not know that *o* is F, or even that *o* exists. For a concrete example of the phenomenon, consider (60) in the following scenario:

Coach: Ann is a six-year-old girl whom Pete, an expert in tennis pedagogy, has never met and whose existence he is unaware of. Pete believes that every six-year-old can learn to play tennis in ten lessons. Jane, Ann’s aunt, is aware of Pete’s feelings on the matter. Jane wants to encourage Ann’s father, Jim, to sign Ann up for tennis lessons, so in conversation with Jim she asserts the following:

(60) Pete believes Ann can learn to play tennis in ten lessons.²²

Ultra-liberal reports have not received much attention in the literature, and to the extent that they have, most all of it has been focused on their unembedded uses. But ultra-liberal attitude reports also have embedded uses. For instance: suppose we don’t know whether Ann is a six-year-old, but we do know that Pete believes all six-year-olds can learn to play tennis in ten lessons. We could then assert:

(61) If Ann is six years old, then Pete believes that she can learn to play tennis in ten lessons.

What is the significance of ultra-liberal attitude reporting as regards third-personal doxastic attitude constructions like (58) and (59), and FP RESTRICTION? Here are two hypotheses about the relationship between ultra-liberalism and restriction: (i) the former is just an instance of the latter; or (ii) the phenomena are fundamentally distinct. If (i) holds, then

²²This example is essentially from (Recanati, 2012, p.152).

given the acceptability of third-personal ultra-liberal reports like (60) and (61), FP RESTRICTION is inadequate as a general constraint on the availability of non-trivial restriction—at least for doxastic attitude constructions. In fact, the acceptability of unembedded ultra-liberal attitude reports like (60) would then be a problem for DYNAMIC RESTRICTION as well. On the other hand, if (ii) holds, constructions such as (58) and (59) pose a problem for FP RESTRICTION only if they involve restriction, and are not instances of (embedded) ultra-liberalism.

Though we lack the space to give the issue the full attention it deserves, we believe there are some *prima facie* considerations in favor of (ii), the thesis that ultra-liberalism and restriction are actually distinct phenomena. One consideration is that an ultra-liberal report like (60) can be acceptable even when it is known to us, for reasons better or worse, that Pete is sure that Ann is a five-year-old (but still holds the belief that all six-year-olds can learn to play tennis in ten lessons). If in circumstances such as these we were to try to restrict Pete’s belief set by the proposition that Ann is a six-year-old, the result would be the empty set. It would then hold trivially that for any proposition, every world in Pete’s restricted belief set would be one in which that proposition is true. As such, we would expect *any* belief ascription with Pete as subject to be acceptable since it will be trivially true, e.g. ‘Pete believes no six-year-old can learn to play tennis in ten lessons’. Clearly this is not what we find.

Another consideration against identifying ultra-liberalism with restriction is that it is significantly harder to find acceptable ultra-liberal reports involving *factive* attitude verbs. Even if we know that Pete knows that no two-year-old can learn to play tennis in ten lessons, unless we have reason to believe that Pete also knows our friend’s two-year-old Little Jimmy’s age, it is hard to hear a true reading of the relevant ultra-liberal knowledge report. By contrast, restriction does not seem to discriminate between factive and non-factive attitude verbs.

So, supposing we take these arguments to show that ultra-liberalism and restriction are indeed distinct phenomena, the argument against FP RESTRICTION from (58) and (59) is sound only if they are instances of the latter rather than the former. But there is a straightforward story on which this is not the case. Let us focus on (58) (‘If Weyland Corporation’s earnings miss expectations, the stock price will fall and our financial advisor is of the opinion that we should buy’). It is natural to suppose that our financial advisor holds a general belief to the effect that one should buy stock in any company that bears the rough financial profile of Weyland Corporation and misses expectations. But since we don’t know whether Weyland Corporation will miss expectations, we can’t yet say that our advisor thinks we should buy. Still, we can say that *if* it misses expectations, then our advisor thinks we should buy. Hence (58).²³ Importantly, however, if we were to discover that Weyland Corporation in

²³The same points apply *mutatis mutandis* to (59). In this case we know Bob holds the general belief that any door

fact missed expectations, we *would* be able to felicitously assert ‘Our financial advisor thinks we should buy’ outright—even if we knew that our financial advisor had no idea whether Weyland Corporation had in fact missed expectations. (58) thus appears to have many of the hallmarks of an embedded ultra-liberal attitude report.

Taking stock

Setting aside these points against echoing uses of attitude constructions and conditionals such as (58) and (59), we do not want to oversell our position here. We are attracted to FP RESTRICTION for a wide variety of attitude verbs, and think there are moves that could be made to explain away some of the recalcitrant uses considered in this section. But we are reticent to claim anything stronger than that.²⁴

In fact, we are now about to argue that there are other attitude constructions that present decisive looking counterexamples to FP RESTRICTION, and thus that it can’t be a general constraint on the availability of restriction. However, we are also about to argue that these attitude constructions present equally decisive looking counterexamples to DYNAMIC RESTRICTION. So although it might appear that this section’s tentative defense of FP RESTRICTION is in vain, there is an interesting backup position that isn’t threatened by the next section’s examples. It is the position that FP RESTRICTION and DYNAMIC RESTRICTION stand and fall together. That is to say, for certain attitude verbs, e.g. emotive factives, non-trivial restriction requires the satisfaction of both FP RESTRICTION and DYNAMIC RESTRICTION, while for other attitude verbs neither is required. It is ultimately out of interest in this thesis that we take seriously the question of whether FP RESTRICTION is robust for the attitude verbs under consideration so far.

bell ringings (in the relevant location/timeframe) will be caused by Steve.

²⁴There is another potential class of counterexamples to FP RESTRICTION, albeit one that requires only a (relatively) superficial revision of the condition. The basic idea is that the proper form of the restriction isn’t to first-person subject terms, but to *interlocutor*-denoting subject terms. Cases like the following are the motivation. Suppose Andrew and Jane are our interlocutors in *Bill’s Holiday*, then (62)-(63) are felicitous:

- (62) [Gesturing towards Andrew] If Bill is on a plane to Cuba, then I bet you are also surprised he left without saying goodbye.
- (63) If Bill is on a plane to Cuba, then I bet Jane also thinks that Bill didn’t tell any of us he would be leaving.

However, barring echoing uses (63) is still unacceptable if Jane is far away and not part of our conversation. If this is correct, then perhaps the relevant constraint should be formulated as follows:

- (64) INTERLOCUTOR-PRESENT RESTRICTION
If Φ is non-trivially restricted by \dagger in ‘S Φ s that P’, then S denotes an *interlocutor* and Φ is in the present tense.

FP RESTRICTION would then just be a special instance of the broader condition on restriction.

6 Some (especially) puzzling cases

This section examines the embedding properties of evidential factives such as ‘know’ and ‘remember’, and the desire verb ‘want’. We argue that despite the discussion in the previous sections, neither DYNAMIC RESTRICTION nor FP RESTRICTION seems to hold for these expressions. The question of what to make of the apparent divergence in the availability of restricted readings between these verbs and the other verbs previously considered is then addressed in §7.

We will take evidential factives to be verbs like ‘know’, ‘remember’, ‘see’, ‘can hear’, and ‘can tell’.²⁵ We believe they display the same sort of puzzling behavior exhibited by the attitude constructions of §1. For instance, consider the following case from [Holguín \(2018\)](#):

Memory Experiment: Joan and Megan are participating in a trial of a drug whose primary effect is to swamp its subject with an extraordinary number of fake “memories” of the events of the past 24 hours. One of the subjects will get the drug, while the other will get a placebo. Who gets which is determined by a coin-flip whose result is known only to the experimenters.

During the experiment Joan and Megan are both (separately) asked ‘Do you remember what you ate for dinner yesterday?’ Joan appears to remember that she ate fish; Megan appears to remember that she ate spaghetti. As a matter of fact it was Joan who got the placebo and Megan who got the drug. Only Joan’s memory is genuine.

Putting ourselves in Joan’s shoes, the following familiar looking attitude constructions are perfectly natural:

- (65) If I got the placebo, then I remember what I ate for dinner last night.
- (66) Either I got the drug, or I remember what I ate for dinner last night.

So to do the following less familiar ones:

- (67) I might remember what I ate for dinner last night; it depends on whether I got the placebo or the drug.
- (68) One of us remembers what she ate for dinner last night.

And, strikingly, knowing that Joan got the placebo, the experimenters may felicitously assert things like:

²⁵Like emotive factives, evidential factives are knowledge entailing (or presupposing). But unlike emotive factives, the extra component is something evidential rather than emotive. Where emotive factives convey the subject’s feelings about the known proposition, evidential factives convey its etiology. We take ‘know’ to be an evidential factive whose evidential component is trivial.

(69) Joan remembers what she ate for dinner last night.

These uses of ‘remember’ seem to be puzzling in just the same way the uses of ‘surprise’, ‘regret’, ‘thinks’, etc., of §1 are. Intuitively Joan doesn’t *know* what she had for dinner last night—at least not in anything like the normal sense of ‘know’—so it’s unclear how she could *remember* it in anything like the normal sense either.²⁶

The same phenomenon arises with ‘know’ itself (this case is also from [Holguín \(2018\)](#)):

History Exam: Peggy and Pete are students in History 101. In preparation for the multiple-choice final exam, both have purchased and subsequently memorized answer sheets from their corrupt teaching assistant Roger. However, moments before the exam Roger shares an unfortunate discovery: due to his ineptitude, one of Peggy or Pete was given answers to an entirely different exam, and he doesn’t know who it was. But he does know that the two answer sheets happen to disagree on the answer to every single question.

Question 5 of the exam is: ‘In what year did the Berlin Wall fall?’ The possible answers are a: 1984, b: 1989, c:1991, and d: 1993. Peggy’s answer sheet says b. Pete’s says c. Peggy is thus the one with good answers.

Putting ourselves in Peggy’s shoes, the following familiar looking attitude constructions are perfectly natural:

(71) If my answer sheet is good, then I know what the answer to question 5 is.

(72) Either my answer sheet is bad, or I know what the answer to question 5 is.

So too are the less familiar ones:

(73) I might know what the answer to question 5 is; it depends on whether I got the good answers.

(74) One of us knows what the answer to question 5 is. (I hope it’s me.)

And supposing Roger the TA eventually discovers that it was Peggy who got the good answers, it is fine for him to assert things like:

(75) Peggy knows what the answer to question 5 is.

²⁶Evidence that ‘remember’ is knowledge entailing is provided by the badness of conjunctions such as (70):

(70) ?? Although I don’t know what I had for dinner last night, I do remember what I had for dinner last night.

We hope the parallels between *Memory Experiment* and *History Exam* are clear. By way of selling the claim that these uses of ‘know’ are non-standard (and thus deserve special semantic treatment), note that, intuitively speaking, Peggy’s credence that the answer to question 5 is b can be no higher than .5. After all, Peggy is well aware that there’s only a .5 chance that she got the good answers, and she’s certain that if she *didn’t* get the good answers then the answer to question 5 isn’t b. To press the point further, notice that we can imagine variants of the case on which Peggy’s credence that the answer to question 5 is b is arbitrarily small. Just imagine, for instance, that there are 1,000 possible answers to a given multiple choice question and that there are 1,000 different students taking the test, each having purchased a different answer sheet from Roger, and that one and only one of the students has the good answers, and that Peggy knows all of this. Still, supposing Peggy is in fact the one with the good answers (and that Roger eventually finds this out), all of (71)–(75) continue to sound fine. So absent a highly non-standard interpretation of the technical term ‘credence’, (71)–(75) present counterexamples to the widely held view that ‘S knows that P’ entails ‘S has a high credence in P’.²⁷ We thus conclude that whatever the semantic contribution of ‘know’ on the natural readings of any of these sentences, it is not plausibly knowledge_T.²⁸

Finally, consider the following case involving ‘want’ adapted from [Jerzak \(forthcoming\)](#):

Wine Time: You’ve been invited to a party and tasked with bringing the wine. You know that the other guests have reasonably strong preferences, but unfortunately you know basically nothing about wine. At the grocery store you have a choice between a Pinot Noir (from California) and a Malbec (from Argentina), which you take to be equally likely to be the optimal choice for the other guests. As such you think to yourself:

(78) If they prefer wine from California, then I want the Pinot Noir.

²⁷See, e.g., [Rothschild & Spectre \(2016\)](#) for some applications of this principle. Note that the principle can also be derived from the claim that ‘S knows that P’ entails ‘S believes that P’, and the claim that ‘S believes that P’ entails ‘S has high credence that P’. These principles are further orthodoxy, though see [Hawthorne et al. \(2016\)](#) for some reasons to doubt the connection between belief and high credence.

²⁸In each of (65)–(69) and (71)–(75), ‘remember’ and ‘know’ take inquisitive complements rather than declarative complements. This is because—for reasons that we lack the space to explore here—analogue constructions with declarative complements tend to sound a bit worse. Still, we will assume that in the context of *Memory Experiment*, ‘Joan remembers what she ate for dinner last night’ entails ‘Joan remembers that she ate fish for dinner last night’, and that in the context *History Exam*, ‘Peggy knows what the answer to question 5 is’ entails ‘Peggy knows that the answer to question 5 is b’. If these entailments didn’t hold, then it would be difficult to explain the abominableness of the following conjunctions:

(76) ?? Joan remembers what she ate for dinner last night, but she doesn’t remember that she ate fish for dinner last night.

(77) ?? Peggy knows what the answer to question 5 is, but she doesn’t know that the answer to question 5 is b.

For further arguments in favor of these entailments, see [Holguín \(2018\)](#).

- (79) Either they prefer wine from Argentina, or I want the Pinot Noir.
- (80) a. I might want the Pinot Noir; it depends on whether they prefer wines from California or Argentina.
- b. There's a chance I want the Pinot Noir; it depends on whether they prefer wines from California or Argentina.

And if an onlooker (somehow) knows of my situation and of the other guests' actual preferences (suppose it's for the Pinot Noir), she may think to herself:

- (81) That person wants to buy the Pinot Noir.

Although 'want' reports are not knowledge entailing, (78)-(81) are still intuitively puzzling. For instance, I have non-trivial credence that the other guests have a preference for the Malbec, so it is difficult to see how (81) could be true. Once again, the attitude verb appears to take a special reading here.

The hypothesis that 'remember', 'know' and 'want' take restricted readings in their respective cases captures the data nicely. Take *History Exam*. Although Peggy does not know_⊤ that the answer to question 5 is b, there is a restriction, G —the set of possible worlds in which Peggy has the good answers—such that she knows _{G} that the answer to question 5 is b. This is because she knows_⊤ that either she has bad answers or the answer to question 5 is b. Thus $(\text{Epi}_{w@, \text{Peggy}} \cap G) \subseteq \{w \mid \text{The answer to question 5 is b at } w\}$. This means we can represent the natural readings of, e.g. (71) and (75) as follows:

- (82) If my answer sheet is good, then I know _{G} what the answer to question 5 is.
- (83) Peggy knows _{G} what the answer to question 5 is.

The same can be said (*mutatis mutandis*) for the uses of 'remember' in *Memory Experiment* and 'want' in *Wine Time*.²⁹ We conclude that these cases provide further support for the bare-boned restriction semantics of §3.

However, it should be clear that some of the cases above present striking counterexamples to both FP RESTRICTION and DYNAMIC RESTRICTION. We will focus on *History Exam* to bring out the relevant points. With regard to FP RESTRICTION, (75) is a straightforward counterexample to the FIRST part, since it is uttered by Roger (the TA) about a third-party (namely Peggy). Counterexamples to the PRESENT part may be generated with ease too:

- (84) If I received the good answers, then I *knew* what the answer to question 5 was.
- (85) If I receive the good answers, then I *will know* what the answer to question 5 is.

²⁹Supposing that we give 'want' a comparative desirability semantics similar to that given to 'regret' in §3.3.

We can imagine Peggy uttering (84) after taking the exam, or (85) prior to having received her answer sheet. In each case the sentence sounds perfectly natural.³⁰

As for DYNAMIC RESTRICTION, all of (73)–(75) pose a problem for this principle, since none involve a conditional, disjunction, or conjunction. Instead, (73) embeds its knowledge report in a modal, (74) embeds it under a quantifier, and (75) is just the unembedded report itself. Indeed, (73)–(75) make clear the difficulty for any condition on restriction in the spirit of DYNAMIC RESTRICTION: the restriction effects witnessed by (73)–(75), if genuine, cannot be derived from any of the formal features of the environments in which the embedded reports occur. Instead, extra-linguistic features of the context must be called upon to supply the relevant restrictions.

Finally, although the main aim of this section has been to show that there are cases that pose a problem for DYNAMIC RESTRICTION and FP RESTRICTION in the sense that these conditions are too strict, it is worth observing that there is also a sense in which these conditions are not strict enough. And that is because it is difficult to get acceptable restricted readings of ‘remember’ and ‘know’ in “tautologous” attitude constructions of the form ‘If P, then I Φ whether P’ and ‘Either P, or I Φ whether not-P’. Consider, e.g.:

- (86) a. ?? If Bill is on a plane to Cuba, then I know whether he is.
b. ?? If Chris injured himself on the tennis court, then I remember whether he did.
c. ?? Either a lot of people are on the deck outside, or I know whether there aren’t.

The badness of these constructions is not predicted by either LOCAL RESTRICTION or FP RESTRICTION. It also marks an interesting contrast between evidential factives and emotive factives: as was pointed out in §2 with (23) (‘If Bill is on a plane to Cuba, then that surprises me’), such conditionals are perfectly felicitous.

However, we should not overstate the extent to which evidential factives are unique here. Though emotive factive tautologous attitude constructions sound fine, doxastic tautologous attitude constructions tend to sound somewhat marginal (even if better than their evidential factive analogs):

- (87) a. ? If it’s raining, then I think it’s raining.
b. ? If Chris injured himself on the tennis court, then I’m sure he did.

These data suggest that the acceptability of tautologous attitude constructions runs on a spectrum: on the one end there are evidential factives, which sound mostly terrible; in the middle there are doxastics, which sound okay but not great (though see Chalmers & Hájek

³⁰Note that the true readings of (84)–(85) are not those on which ‘know’ is forced to take its normal interpretation. Peggy can utter (84) knowing_T full well that her past credence that the answer to question 5 is b was no greater than .5.

(2007) for arguments that take the acceptability of these constructions for granted); and on the other end there are emotive factives, which sound quite natural. What would explain this distribution is a question for further research.³¹

7 Looking ahead

7.1 A semantic cleft?

What are we to do with these observations about evidential factives and ‘want’? We see two ways of going forward, broadly construed. The first way is to take the data just surveyed to be suggestive of a semantic cleft between different categories of attitude verbs. This view grants that all attitude verbs take non-trivially restricted readings, but denies that the underlying mechanisms responsible for such readings is the same in each case. In favor of this view one could cite the fact that restricting attention to, e.g. emotive factives, DYNAMIC RESTRICTION and FP RESTRICTION both look quite robust (albeit the former more so than the latter); by contrast, restricting attention to, e.g. evidential factives, DYNAMIC RESTRICTION and FP RESTRICTION both seem straightforwardly problematic.

We see this as an intriguing place for further research. But we note that it is not without obvious theoretical drawbacks. Most plainly, it would be surprising if the true account of verb restriction requires two independent mechanisms that just happen to generate similar patterns of readings across categories of attitude verbs. Also notice that no matter what account is given of, e.g. emotive factives, the data surveyed in §6 suggests that the constraints on, e.g. ‘know’ will have to be considerably complex. So, one will not be able to avoid positing theoretically unwieldy restriction mechanisms just by countenancing a semantic cleft.³²

³¹Note too that any explanation of these data will have to reckon with the fact that judgments here are extremely sensitive, in the following sense: evidential and doxastic tautologous attitude constructions improve dramatically when of the form ‘If P, then I Φ that not-not-P’, ‘If P, then I Φ that P or (\neg P and Q)’, etc.

³²Also note that the cleft will probably not be in keeping with existing taxonomies of attitude verbs. For example: ‘hope’ seems to obey DYNAMIC and FP RESTRICTION, despite the fact that ‘want’ doesn’t. To see this, suppose you’ve already bought the bottle (either the Pinot Noir or the Malbec) but you don’t know which one you bought. In this setting the following thoughts seem natural:

- (88) a. If the guests prefer wine from California, then I hope I bought the Pinot Noir.
b. Either the guests prefer wine from Argentina, or I hope I bought the Pinot Noir.

So, ‘hope’ exhibits restriction effects (we already saw evidence of this in §1). But in keeping with DYNAMIC and FP RESTRICTION, it can only do so in the right linguistic environments:

- (89) a. ?? There’s a chance that I hope I got the Pinot Noir; it depends on whether they prefer wine from California or Argentina.
b. ?? [Uttered by an onlooker:] That person hopes they bought the Pinot Noir.

Alternatively, one might prefer to “generalize to the worst case” and posit a single mechanism that accounts for the availability of restricted readings of all attitude verbs. On an approach of this sort, one would enrich the semantics for attitude verbs with a parameter for restriction—just in the way of §3’s entries for ‘think’, ‘know’, and ‘regret’—but more or less leave it up to context to supply its value. The task of accounting for the general (even if not entirely universal) patterns of DYNAMIC RESTRICTION and FP RESTRICTION would then have to be treated separately.³³

Overall, what seems clear to us is that a view comprised of DYNAMIC RESTRICTION and FP RESTRICTION has a systematicity and elegance to it that suggests it is worthy of further investigation; but also that there is enough recalcitrant data that other, less constrained alternatives should be on the table too. As such, our goal has simply been to lay out some of the options in the hopes of stimulating further research.

7.2 Other issues

We end this section by discussing two sets of issues naturally raised in light of the central claims of the paper.

First, many theorists have maintained that *epistemic modals* exhibit restricted readings.³⁴ This naturally raises the question of the relationship between the mechanism that generates restricted readings of verbs and the mechanism that generates restricted readings of modals. A reasonable starting hypothesis is that it is the same underlying mechanism, at least at some

-
- c. ?? [Uttered by an onlooker:] If the guests prefer wine from California, then that person hopes they bought the Pinot Noir.

Given the similarities between ‘hope’ and ‘want’, these observations come as something of a surprise.

³³For reasons whose proper treatment is beyond the scope of the present discussion, we believe that if DYNAMIC RESTRICTION is abandoned, then the object that constrains the denotation of attitude verbs should be derived from a *partition* of logical space, i.e. from the semantic value of a *question*. The restriction effect would then be achieved by intersecting, e.g. $Epi_{w,S}$ (in the case of ‘know’) with the proposition that is the true answer to that question at the relevant point of evaluation. (This will normally be the actual world, but on certain embedded uses—i.e. modals and conditionals—the point can shift to various non-actual worlds.) The move to higher type is needed to account for the true readings of ascriptions like:

- (90) There is a 50% chance that I know what the answer to question 5 is.

If ‘know’ denotes $knowledge_G$, then (90) is false. After all, Peggy *knows*_⊥ that she *knows*_G what the answer to question 5 is. However, if ‘know’ denotes $knowledge_{G?}$ where $G?$ is a question along the lines of ‘Do I have the good answers?’, then we get the reading we want. If Peggy’s answer sheet is good, then the restriction is the proposition that her answers are good, and on that restriction she knows the answer to question 5. But if Peggy’s answer sheet is bad, then the restriction is the proposition that her answers are not good, and on that restriction she *doesn’t* know the answer to question 5. There’s only a 50% chance her answers are good. Hence, there’s only a 50% chance the restriction is the proposition that her answers are good. Thus, there’s only a 50% chance she *knows*_{G?} what the answer to question 5 is.

³⁴Some version of this idea is just about universally accepted in the literature on epistemic modals. See, e.g., (Yalcin, 2007), (Gillies, 2010), (Kratzer, 2012), (Dorr & Hawthorne, 2013), and (Moss, 2015) for a sampling of views that diverge on a number of matters on the semantics of epistemic modality, but not on the fact that modals are sometimes subject to restriction-like effects.

level of abstraction. But supposing this hypothesis is correct, a number of questions present themselves. For starters, what are we to make of Yalcin's (2007) famous minimal pair?:

- (91) a. ?? If it is raining and it might not be raining, then...
b. If it is raining and I don't know that it is raining, then...

In particular, should we expect there to be infelicitous (restricted) readings of sentences like (91b)?³⁵ Additionally, if we are attracted to the view that *doesn't* posit a semantic cleft between the verbs that obey DYNAMIC and FP RESTRICTION and those that don't, then it looks like the workings of restriction cannot in general be explained in terms of the formal properties of the constructions in which restrictable expressions occur. Is that some reason to prefer theories of epistemic modals that explain their restriction effects in terms that are fundamentally extra-linguistic?³⁶

Second, there are important issues about the kinds of *mental states* attitude verbs denote on their restricted readings. For instance: supposing the view sketched in §3 is correct and we traffic in relations such as knowledge_† in ordinary discourse, then we have reason to believe that 'know' sometimes denotes relations that are in gross violation of, e.g., credence constraints on knowledge. What, then, are the epistemological consequences of the availability of these relations in our talk of knowledge? Do cases like *Memory Experiment* and *History Exam* present putative counterexamples to much of the orthodoxy about knowledge, or are certain resolutions of † privileged for the purposes of epistemological theorizing?³⁷ These questions deserve further investigation.

8 Conclusion

Our paper opened with a puzzle: certain ordinary uses of attitude verbs present *prima facie* counterexamples to standard logical inference rules (e.g., modus ponens and disjunctive syllogism). We resolved the puzzle by positing restricted readings of these verbs. We then turned to the task of finding mechanisms that would explain when the restriction arises and

³⁵There is some evidence that we should. Consider:

- (92) ? If Bill is in Cuba right now but I don't know whether he is sleeping at home, then...
(93) ? If Federer will win Wimbledon but I don't know whether he will win any slams, then...

To our ears these conditionals sound quite bad. Putting things rather roughly, they have the phenomenology as of being asked to suppose something and then to forget that one is supposing it. The hypothesis that in each case 'know' takes a restricted reading would, if true, be a nice explanation of these facts.

³⁶The similarities between the attitude construction data of §6 and some of the embedded epistemic modal data of Dorr & Hawthorne (2013) and Moss (2015) could be taken as evidence in favor of a positive answer to this question.

³⁷See Holguín (2018) for extended discussion of the epistemological issues raised by a semantics for 'know' along the lines presented in §3.

what its content is when it does. Our first attempt connected restriction to the dynamic effects of the connectives (codified in *DYNAMIC RESTRICTION*). The resulting story had simple, clean predictions: attitude verbs, when restricted, are always restricted by dynamically supplied propositions. When an attitude verb appears in the consequent of a conditional, the underlying set of possibilities is (optionally) restricted by the content of the conditional's antecedent; when an attitude verb appears in the second disjunct of a disjunction, the underlying set of possibilities is (optionally) restricted by the content of the negation of the first disjunct; etc.; otherwise the verb always takes its normal reading. We then observed that the presence of restriction also seemed to depend on the verb's subject being in the first person and its tense being present, and so added *FP RESTRICTION* as a further constraint (with reservations). But then evidential factives and 'want' entered the picture, and both *DYNAMIC* and *FP RESTRICTION* were found to be wanting—at least on the assumption that the underlying phenomenon is a unified one.

Although we believe the entries for verbs like 'think', 'know', and 'regret' stated in §3 are ultimately correct, we are without a general account of the underlying mechanisms that tells us when a restriction will appear and what its content will be when it does. We believe the task of finding such an account will prove to be highly non-trivial. But we also believe it should be a matter of broad theoretical interest. For not only is the question of the proper semantics for attitude verbs interesting in its own right, but given the nature of the data we are dealing with, there are reasons to believe that its investigation may shed new light on debates about epistemic modals and the nature of knowledge.

References

- Beaver, David I. 2001. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications.
- Bennett, Jonathan. 2003. *A Philosophical Guide to Conditionals*. Oxford University Press.
- Blumberg, Kyle, & Holguín, Ben. 2018. Ultra-Liberal Attitude Reports. *Philosophical Studies*, **175**(8), 2043–2062.
- Brogaard, Berit, & Salerno, Joe. 2013. Fitch's Paradox of Knowability. In: Zalta, Edward N. (ed), *The Stanford Encyclopedia of Philosophy*, winter 2013 edn. Metaphysics Research Lab, Stanford University.
- Chalmers, David J., & Hájek, Alan. 2007. Ramsey + Moore = God. *Analysis*, **67**(294), 170–172.
- Dorr, Cian, & Hawthorne, John. 2013. Embedding Epistemic Modals. *Mind*, **122**(488), 867–914.
- Drucker, Daniel. 2017. Policy Externalism. *Philosophy and Phenomenological Research*, **94**(3).
- Egré, Paul. 2008. Question-Embedding and Factivity. *Grazer Philosophische Studien*, **77**(1), 85–125.
- Giannakidou, Anastasia. 2006. Only, Emotive Factive Verbs, and the Dual Nature of Polarity Dependency. *Language*, **82**(3), 575–603.
- Gillies, Anthony. 2010. Iffiness. *Semantics and Pragmatics*, **3**(4), 1–42.
- Hawthorne, John, Rothschild, Daniel, & Spectre, Levi. 2016. Belief is weak. *Philosophical Studies*, **173**(5), 1393–1404.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, UMass Amherst.
- Heim, Irene. 1992. Presupposition Projection and the Semantics of Attitude Verbs. *Journal of Semantics*, **9**(3), 183–221.
- Hintikka, Jaakko. 1962. *Knowledge and Belief*. Ithaca: Cornell University Press.
- Holguín, Ben. 2018. *Knowledge by Constraint*. Unpublished manuscript.
- Jerzak, Ethan. forthcoming. Two Ways to Want. *Journal of Philosophy*.
- Karttunen, Lauri. 1974. Presupposition and Linguistic Context. *Theoretical Linguistics*, **34**(1), 181–.
- Klinedinst, Nathan, & Rothschild, Daniel. 2012. Connectives Without Truth Tables. *Natural Language Semantics*, **20**(2), 137–175.
- Kratzer, Angelika. 1986. Conditionals. *Chicago Linguistics Society*, **22**(2), 1–15.
- Kratzer, Angelika. 2012. *Modals and Conditionals*. Oxford University Press.
- Levinson, Dmitry. 2003. Probabilistic Model-theoretic Semantics for 'want'. *Semantics and Linguistic Theory*, **13**(0), 222–239.
- Lewis, David K. 1973. *Counterfactuals*. Blackwell.
- May, Robert. 1985. *Logical Form: Its Structure and Derivation*. MIT Press.
- Moss, Sarah. 2015. On the Semantics and Pragmatics of Epistemic Vocabulary. *Semantics and Pragmatics*, **5**(8), 1–81.

- Pryor, Jim. 2004. An Epistemic Theory of Acquaintance. *In: UCSB Conference on Content and Concepts*. 2004.
- Recanati, François. 2012. *Mental Files*. Oxford University Press.
- Rothschild, Daniel, & Spectre, Levi. 2016. A Puzzle About Knowing Conditionals. *Noûs*.
- Schaffer, Jonathan, & Szabo, Zoltan Gendler. 2013. Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions. *Philosophical Studies*, 1–53.
- Schlenker, Philippe. 2009. Local Contexts. *Semantics and Pragmatics*, 2(3), 1–78.
- Stalnaker, Robert. 1974. Pragmatic Presuppositions. *Pages 47–62 of: Stalnaker, Robert (ed), Context and Content*. Oxford University Press.
- von Stechow, Kai. 1999. NPI Licensing, Strawson Entailment, and Context Dependency. *Journal of Semantics*, 16(2), 97–148.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- Yalcin, Seth. 2007. Epistemic Modals. *Mind*, 116(464), 983–1026.
- Yalcin, Seth. 2012. A Counterexample to Modus Tollens. *Journal of Philosophical Logic*, 41(6), 1001–1024.