

Attitudes, Conditional and General*

Daniel Drucker

1 The Problem of Non-Doxastic Attitude Conditionals in Ignorance Contexts

Semanticists use intuitions about the truth of utterances in real or imagined contexts to constrain their explanations of the meanings of expressions. Even though that's the right thing to do, it introduces a potential problem: sometimes ordinary speakers, and even sometimes trained semanticists, fail to see a *way* in which an utterance can be true. This can happen through, for example, inattention to some kind of context sensitivity, as when people wonder whether 'that elephant is tiny' can ever be truly uttered. It can also happen because paradigms (or "prototypes") confuse people about possible extensions: even quite liberal people might forget that, at least from their perspective, 'this man is pregnant' can be true. By reminding them that you might be talking about how a brontosaurus would fare in a fight with the elephant or of their general social views, the intuition that the utterance can in fact be true will follow. But sometimes we can see *that* an utterance is true, but it's hard to see how it could be true. This can lead to revisionary semantic accounts of some of the expressions involved; but some of these will be unnecessary because there is in fact a straightforward way in which the utterances can be true, a way that requires saying very little new or special. I'll argue this kind of trap

*This paper has been a long time in production. Thanks especially to Eric Swanson, with whom I discussed most of these ideas in the early stages and to Dmitri Gallow, who read many different portions of it over the years. Thanks also to Kyle Blumberg, Josh Dever, Ben Holguín, Harvey Lederman, Matt Mandelkern, Henry Schiller, and participants in a Covid Zoom philosophy of language workshop. Thanks also to Hakeem Jefferson for providing an excellent work environment, and to Oscar Zoletto. Finally, extra special thanks to three reviewers for this journal and the editor Paula Menéndez Benito for extremely careful help.

is well worth avoiding when it comes to the kind of expression I'm writing about here.

Following Blumberg and Holguín (2019), I'll call any utterance of the form \ulcorner if ϕ , X \lrcorner that ψ the utterance of an *attitude conditional*.¹ Some of them are not mysterious:

- (1) If Jeffrey hasn't found a partner yet, then he's still sad that he's single.

Suppose he hasn't yet found a partner. Since he then knows he hasn't, we know he'd be sad about it. And indeed that knowledge will explain, probably causally explain, his sadness that he's still single. This contrasts with the following case, where the speaker knows that while they have yet to hear from an academic job they applied to, other applicants have had interviews arranged. Now consider:

- (2) If my application for the position really was rejected, I'm very sad that it was.

In this case, the speaker likely *knows* that they don't know whether their application was in fact rejected. And yet the conditional still seems like it very well might be true. The challenge is to say how it might be. Let me first be a little more precise about the target phenomenon, and then a little more explicit about the challenge.

I'll say that an attitude conditional of the form \ulcorner if ϕ , I \mathcal{A} that ψ \lrcorner is uttered in an *ignorance context* if and only if the speaker doesn't know whether ϕ . For my purposes, then, ignorance contexts will only apply to first-personal attitude conditionals. Sadness, considered as a propositional attitude, is *non-doxastic*, meaning, roughly, that it's a propositional attitude that isn't belief, knowledge, or credence. I say "roughly" because there are tougher cases, including being surprised and being disappointed. In this paper, I will focus on non-doxastic attitude conditionals, and I'll err on the side of caution and discuss only those attitudes that are *clearly* non-doxastic. Toward the end, though, I will return to doxastic or borderline-doxastic attitudes, since they provide an important

¹Here, very briefly, are some of the notational conventions I'll be using. First, corner, or Quine, quotation marks ("for all sentences ϕ , $\ulcorner\phi \vee \neg\phi\lrcorner$ is true") are kinds of quotation mark that allow variables within the quotation to range over expressions; the example I just gave says that every sentence that disjoins a sentence with its negation is true. (In practice the universal quantification will usually be implicit.) Greek variables in the ϕ , ψ , χ sequence, along with their primed or subscripted counterparts, stand for sentences of ordinary English. ' \mathcal{A} ' is a distinguished predicate variable ranging only over attitude expressions ('hope', 'regret', etc.). Lower-case Roman letters in the p , q , r sequence are variables ranging over propositions. I'll introduce more notation in the text as it becomes helpful.

challenge for future work to address, for me and for others as well. So, in addition to (2), here are some of the other examples I'll be working with:

- (3) If my application for the position really was rejected, I'm glad that I won't have to live in Whoville.
- (4) If my application for the position really was rejected, I regret that I wasted the time applying to such a longshot.
- (5) If my application for the position really was rejected, I'm angry that they gave me so much false hope.
- (6) If my application for the position really was rejected, I hope that they tell me I didn't get the position soon.

It's easy to come up with more examples once you understand the basic structure.

The challenge, again, is to explain how non-doxastic attitude conditionals like (2) through (6) might be true in their natural ignorance contexts. I've already suggested a bit about why the challenge is a tough one, but let's get into some details.

I will assume the following about all the attitude conditionals: if \ulcorner if ϕ , I \mathcal{A} that $\psi\urcorner$ is true in c and $\ulcorner\phi\urcorner$ is also true in c , then \ulcorner I \mathcal{A} that $\psi\urcorner$ is also true in c . This is a simple application of a semantic version of *modus ponens*. For ease of reference, I'll call the speaker in (2) through (6) Todd. Next, suppose the antecedent of each of these conditionals is true, i.e., that Todd's application really was rejected. Then, if (2) is true in c , then it's also true in c that Todd is very sad that his application was rejected. But propositional sadness, as well as gladness, regret, and anger, is often taken to be *knowledge-entailing*, i.e., if p is the content of S 's sadness, then S knows that p .² In fact all the attitudes expressed by what Anand and Hacquard (2013) call "emotive factives" at least look knowledge-entailing. Examples like the following constitute the main evidence:

²See, e.g., Unger (1975), Gordon (1987), Williamson (2000), and most recently Dietz (2018). I formulated this property so as to be neutral about whether it would follow from the nature of the attitudinal states (sadness, etc.) or whether it's a conventional semantic fact about attitude verbs, for example that they are lexically generated presuppositions.

- (7) ???I have no idea whether my application for the position really was rejected, but I'm very sad that it was.
- (8) ???Todd has no idea whether his application for the position really was rejected, but he's very sad that it was.

Examples like (7) and (8) sound horrible; they seem to be a more general version of what is sometimes called “Moore’s paradox”, that it sounds somehow very bad to utter sentences like $\lceil \phi, \text{ but I don't believe that } \phi \rceil$ or even $\lceil \phi, \text{ but I don't know that/whether } \phi \rceil$. So far, that’s just data; it stands in need of explanation. With (7) and (8), that sadness is knowledge-entailing is a natural explanation. In that case, we would have a problem: (2) would then entail in *c* that Todd *knows* that his application was rejected. But that contradicts our assumption that this was an ignorance context. So, we either have to deny that sadness really is knowledge-entailing, or we have to have some other explanation of examples like (7) and (8) that don’t assume that it is knowledge-entailing. Call this *the problem of unwarranted detachments*.

Some philosophers use Moore-paradoxical utterances to motivate “norms of assertion”, rules speakers are somehow on the hook for following in making assertions. The most famous such norm is:

KNOWLEDGE NORM OF ASSERTION. *S* must: assert that ϕ only if *S* knows that ϕ .³

Perhaps, then, the issue with (7) and (8) is that although the consequent can be *true* in that context, it can’t be known, and thus it can’t be felicitously asserted because of the knowledge norm. Notice also that some weaker norms than the knowledge norm, for example one where we replace “*S* knows that” with “is reasonable for *S* to believe”, could lead to the same place, since it won’t be reasonable for Todd to *believe* that he’s very sad that he was rejected.⁴ We would still be on the hook for explaining why he can be very sad about it without knowing that it happened, but it still feels like progress.

(8) would still present a problem. In an ignorance context, the speaker—Todd—doesn’t know that his application was rejected. But that doesn’t prevent a third party

³See especially Unger (1975) and Williamson (1996).

⁴For the reasonable-to-believe norm, see Lackey (2007), and for a related norm see Douven (2006).

from knowing it. So, suppose a third party, Diane, overhears Todd saying (2). In that case, Diane can still detach the consequent, using *modus ponens*, inferring that Todd is very sad that his application was rejected. If doing this preserves the Diane’s knowledge, as it seems to me it very well might and typically would do, then she should be in a position to assert (8). And yet it sounds infelicitous. This assumption that the *modus ponens* inference can preserve the third party’s knowledge will follow from what epistemologists call a *closure* condition on knowledge, which I put in a footnote. I also just want to stress that that such an inference can and often does preserve knowledge is itself more plausible than theoretical principles like closure.⁵ So, it seems, the knowledge norm and related norms don’t seem like they’ll help with the problem of unwarranted detachments.

That problem afflicts attitudes expressed by emotive factives, as I said. ‘Hope’, which figures into (6), is not such a verb; hoping that ϕ obviously doesn’t entail that ϕ . If anything, the opposite is true: it’s more plausible that if $\ulcorner S$ hopes that $\phi \urcorner$ is true in c , then $\ulcorner S$ knows that $\phi \urcorner$ is *false* in c .⁷ So perhaps the problem of unwarranted detachments doesn’t afflict ‘hope’. Still, a different problem does. Consider this elaboration of (6):

- (9) If my application for the position really was rejected, I hope that they tell me that I didn’t get the position soon, because sitting here with false hope is awful.

But now consider only the consequent of (9):

- (10) I hope that they tell me that I didn’t get the position soon because sitting here with false hope is awful.

If Diane could detach (10) when (9)’s antecedent is true, then Diane’s in a position to know that (9) is true of Todd. But (10) seems to presuppose that Todd, right then,

⁵I prefer formulations of closure like these:

EPISTEMIC CLOSURE. If S knows that ϕ , and competently deduces p , thereby coming to believe p while retaining the knowledge that ϕ throughout, then S is in a position to know p .⁶

But as I said, the problem doesn’t require anything like a generally true closure principle.

⁷See Heim (1992) for discussion of related inference involving ‘desires’, such that (utterances of) $\ulcorner S$ desires that $\phi \urcorner$ presupposes that S does not believe that ϕ or that $\neg\phi$. To what extent ‘desire’-ascriptions really license these inferences is hard to say; it’s an old issue dating to the *Symposium*, and discussed in Kenny (2003, pages 115–116) and Matthews and Cohen (1967). Heim’s apparent counterexamples to the inference (e.g., (her 41)) ‘(John hired a babysitter because) he wants to go to the movies tonight’ and (her 42) ‘I want this weekend to last forever. (But I know, of course, that it will be over in a few hours.)’ are neither one easily adapted to counterexamples for ‘hope’.

has *false* hope, i.e., that his application has already been rejected. More generally, the following is at least plausible on one disambiguation of ‘because’: if $\ulcorner S \mathcal{A}$ s that ϕ because $\psi \urcorner$ is true in c , then $\ulcorner S$ believes that $\psi \urcorner$ is true in c .⁸

The disambiguation I have in mind is when ‘because’ ascribes the *reasons for which* the ascriber is said to do what they do. Notice that some utterances of $\ulcorner S \mathcal{A}$ s because $\psi \urcorner$ are paraphrasable with $\ulcorner S$ ’s reason for \mathcal{A} ing is that $\psi \urcorner$. Thus:

(11) I went to the store because I wanted some apples. \Leftrightarrow My reason for going to the store was that I wanted some apples.

(12) I stepped on my neighbor’s foot because the bus jolted forward. \nLeftrightarrow My reason for stepping on neighbor’s foot was that the bus jolted forward.

On the typical interpretations of their expressions, the left- and right-hand-sides of (11) necessarily have the same truth-value, whereas the left- and right-hand-sides of (12) may easily differ in truth-value. The entailment I’m describing as generated by (10), call it a *reasons entailment* for a shorthand, only applies to ‘because’ in (11) rather than (12).

The philosophical motivation for the reasons entailment is that something can’t be someone’s reason for forming an attitude if they don’t *believe* it. It does have linguistic indicia. Consider:⁹

(13) ??I had no idea I wanted apples, but I went to the store because I wanted apples.

That utterances like (13) sound infelicitous is some good linguistic evidence for the reasons entailment, even beyond its philosophical plausibility.

Given this entailment, Diane would be in a position to infer Todd believes his hope is false, i.e., that he won’t get the position. But we can simply stipulate that, because Todd knows he wouldn’t know his application has been rejected if he were to believe it, he doesn’t even believe it. Even then, he can seemingly truly utter (9). To once again appeal

⁸See, e.g., Hyman (1999), Alvarez (2016), and Littlejohn (2012), among many others; many authors in this literature would replace ‘believes’ with ‘knows’, but that’s entailed by my version and the difference won’t make a difference for my purpose.

⁹Be sure to interpret ‘because’ as in (11). There’s a way (13) can be true if somehow my desire causes me to go to the store without being my reason. But this is a different use of ‘because’.

to Anand and Hacquard (2013)’s helpful taxonomy, even if “emotive doxastics” (‘hope’, ‘fear’, etc.) and “desideratives” (‘desire’, etc.) don’t have the problem of unwarranted detachments, they have this new problem, call it *the problem of absent beliefs*.

If we want to affirm the possible truth of attitude conditionals like (2)–(6) in ignorance contexts, we need to solve the problems of unwarranted detachments and absent beliefs. As I see it, we have these options:

- *Deny that these attitude conditionals are really conditionals with the basic logical form they seem to have.* According to this strategy, we would say that although (2)–(6) look like real conditionals with a simple, face-value logical form, they’re not. The most obvious version of this strategy is to say that really the attitude verb takes wide scope over the real conditional. Hence (2) is really more like:

(14) I’m very sad that, if my application for the position really was rejected, it was [rejected].

That is, of course, a terrible paraphrase of (2), since Todd isn’t sad about a tautology.¹⁰ Wide-scoping strategies can of course get more sophisticated than this, but this is a basic enough worry to make us properly pessimistic about them.

On a different approach, we posit implicit context sensitivity not unlike similar approaches to conditionals with modals in their consequents:

(15) I prefer candidate A to B and B to C, so if indeed A is no longer viable, I ought to support B in the upcoming election.

In this example, it’s natural to think that, somehow, ‘ought’ introduces quantification over a contextually given domain, which is then restricted by the information expressed by the antecedent.¹¹ Returning now to attitude verbs, the detachments wouldn’t be unwarranted so much as misleading; with (2), for example, we can infer something that has as a more explicit paraphrase something like ‘assuming my application was rejected, I’m very sad that it was’; without the ‘assuming’ bit,

¹⁰Drucker (2019) makes this point.

¹¹Kratzer (1977) is of course the *locus classicus* for this view of the context sensitivity of modals.

hearers would be easily fooled into thinking the person is sad (more) unrestrictedly about it.¹² There's nothing as strikingly wrong with this view as there was with the wide-scope view.¹³ Still, even the following tends to sound bad:

(16) Assuming his application for the position was rejected, Todd is sad that it was.

Even this can sound pretty bad to say. For that reason, Blumberg and Holguín add the following condition:

FIRST-PERSONAL RESTRICTION. If \mathcal{A} is non-trivially restricted by the proposition expressed by $\ulcorner \phi \urcorner$ in $\ulcorner S \mathcal{A} \urcorner$ that ϕ , then S is a first-personal pronoun and \mathcal{A} is in the present tense.

This restriction is necessary to account for why, even though (16) would otherwise be entailed in the given context by (2), it nevertheless cannot express what the consequent of (2) in the first-person can express, given the FIRST-PERSONAL RESTRICTION. (If it did express the same proposition, it would have to be true.)

That said, I have three objections to this approach:

1. It's unnecessary given the data I will introduce in section 2; in fact, I'll argue there that we're committed no matter what to the unrestricted detachments. As a strategy for avoiding the problem of unwarranted detachments, then, it doesn't work, since we can't actually avoid those detachments.
2. The context sensitivity seems not to help with the problem of absent beliefs; even if the *hope* a person is committed to by (9) is only one that would be reported by 'assuming my application was rejected, I hope...', we nevertheless end up inferring that the agent has the *belief* that it was rejected, and the

¹²This is Blumberg and Holguín (2019)'s approach; Silk (2014b) is a structurally similar account of similar detachment issues involving deontic modals.

¹³There's more that's important and insightful about their view than I can go into here or in the paper more generally. For example, their account extends to "attitude disjunctions" like 'either my application was accepted, or I'm very sad that it wasn't'. I think what I say will extend straightforwardly to such disjunctions but I don't pursue that here.

belief seems like it'd be unrestricted. After all, there's only one attitude verb in the sentence, 'hope' rather than 'believes'. Dialectically, this is the less important objection for me; it's only a challenge to the kind of contextualist I'm describing to explain what's going on in this case.

3. FIRST-PERSONAL RESTRICTION feels arbitrary and unexpected given the modal inspiration of the view. 'Assuming' and 'in view of' seem always to successfully restrict the domain of worlds modal expressions quantify over. Why, exactly, should such restrictions not "work" when otherwise the attitude expressions would behave exactly like modal expressions? In section 3, I will argue directly that FIRST-PERSONAL RESTRICTION is not true, and so we don't have to take that explanatory burden on.

I'll present one further objection to this kind of account at the end of the section.

- *Reject modus ponens's unrestricted validity in application to natural language.* There are of course existing alleged counterexamples to *modus ponens*; it's not sacrosanct. McGee (1985) has the following kind of case:

(17) If a Republican will win the 1980 election, then if Reagan doesn't win, then Anderson will win.

Suppose a Republican will win. Then we can infer that if Reagan doesn't win the election, Anderson will win. But this conditional seems false. There are many responses to this possible counterexample, but I hope it's clear nothing in attitude conditionals is like it. The *Miners' puzzle* gives us a more similar case. Here's a quick sketch.¹⁴ Suppose there are two shafts, A and B, and ten miners trapped either in A or B—we have equal credence as to which. We can block water from flowing into A, block it from flowing into B, or block neither. If they're in A (or B) when we block just A (or B), they all live; if they're in A (or B) when we block just B (or A), they all die; and if we direct the water to both A and B, all but one

¹⁴See Parfit (1984) and Regan (1980) for the original puzzle.

of them will survive. People typically judge the following utterances about the case felicitous:

- (18) a. If they're in A, we ought to block A.
 b. If they're in B, we ought to block B.
 c. We ought to flood neither A nor B.

But, at least in the presence of *modus ponens*, these three lead to a contradiction (assuming no relevant equivocation).¹⁵ As I said, unlike the McGee examples, this is closer to our present cases. We have, after all:¹⁶

- (19) a. If they're in A, I want to block A.
 b. If they're in B, I want to block B.
 c. Because I have no idea where they are, I want to block neither A nor B.

(a) and (b) contradict (c) in the presence of *modus ponens*. In fact, the problems are so similar that I take them to be the same problem. Even though some do use structures like (18) and (19) to cast doubt on *modus ponens*, others do not,¹⁷ and anyway we should aim to preserve it if we can, given how strongly intuitive it is. My account will preserve it. It will turn out that we otherwise wouldn't just have to give up *modus ponens* but also universal elimination (i.e., if $\ulcorner (\forall x)Fx \urcorner$ is true in c , then $\ulcorner Fa \urcorner$ is true in c , for any referring name $\ulcorner a \urcorner$).

- *Give a principled reason for denying knowledge-entailingness and the reasons entailment.* We are left with the last option, to give principled reasons why sadness and the rest are not knowledge-entailing and why the reasons entailment fails. That is what I will do in the next section: I will argue that we should turn our attention to other constructions, non-doxastic attitude expressions with 'ever' free relative complements (\ulcorner whatever is $F \urcorner$, etc.). These constructions figure into unproblematic

¹⁵See Kolodny and MacFarlane (2010). Here's the reasoning. Either they're in A or they're in B. Suppose they're in A. Then we ought to block A, so it's not true that we ought to block neither. Now suppose they're in B. Then we ought to block B, so it's not true that we ought to block neither. Either way, it's not true we ought to block neither. But that contradicts (c).

¹⁶Drucker (2019) draws a similar connection between attitude conditionals and the Miners' puzzle.

¹⁷See, e.g., Cariani et al. (2013), Silk (2014a), and von Fintel (2012), among many others.

attitude ascriptions that nevertheless entail the non-doxastic attitude conditionals, even in ignorance contexts. As I said, this makes contextualist accounts of these constructions unnecessary. These ascriptions show why these attitudes aren't knowledge-entailing, and also show why the reasons entailment is also false.

My account will then be a version of the third strategy, which is conservative at the level of semantics but perhaps philosophically more revisionary.

As I've already made clear, this isn't the first paper to discuss these conditionals. To my knowledge, Drucker (2019) was the first to discuss them as *prima facie* surprising and as potential counterexamples to *modus ponens*.¹⁸ He did not give an account of what explains how the conditionals could be true in anything like full semantic detail. Jerzak (2019) discusses 'want'-conditionals in attitude contexts, specifically examples like:

(20) If my comrades would most enjoy the merlot, I want to get the merlot.

To account for them, he develops a MacFarlane (2014)-style relativistic semantics. The idea should be familiar from discussions of, for example, predicates of personal taste like 'tasty', as well as with information-sensitive expressions like 'ought'. First, we enrich our indices with an information state parameter tied to the context of assessment. Second, we assume that a conditional $\lceil \text{if } \phi, \psi \rceil$ is true relative to a world w and information state i just in case $\lceil \psi \rceil$ is true relative to w and i intersected with the content expressed by $\lceil \phi \rceil$. Drucker (2019) points out that without special stipulation, an analogous account of 'am sad', etc. will make the following come out true and assertible in our context, so long as (2) is known to be true and his application was (unbeknownst to him) rejected:

(21) Todd is very sad that his application was rejected.

Of course, maybe we can convince ourselves to be comfortable with that; but if so, that work needs to be undertaken, and that's a large part of what I'll do in the paper. His discussion, unlike mine, is limited to 'want' rather than the full range of non-doxastic attitude expressions that exhibit similar behavior; and like the contextualism I described

¹⁸Iatridou (1991, pages 26–29) discusses examples I would call attitude conditionals, but she appeals to them to argue for movement constraints and obligations in conditionals.

above, my account will undercut the need for his. I'll discuss his data, and how I can account for it without any novelty in the semantics in section 3. Blumberg and Holguín have a stipulation in their semantics requiring the conditionals be in the first person. Again, this is the kind of thing we should try to explain before resorting to stipulation.

In many ways, von Fintel and Pasternak (n.d.)'s account of attitude conditionals is closest to the one I'll present here. (It, too, then qualifies as a face-value account.) They are focused on the following sentence:

(22) If Laura becomes a zombie, she wants you to shoot her.

The account they give aims to derive a good reading for the sentence, one that would be intelligibly felicitous in the relevant circumstances, that doesn't go beyond ordinary resources for understanding conditionals. They couch things in the language of situation semantics, and make the desire expressed in (22) the desire *of* a non-actual situation in which she becomes a zombie that the addressee shoot her in that world. They say: "We thus predict [(22)] to denote a proposition true of a situation s iff in all of the most likely worlds w from s in which Laura becomes a zombie, the minimal zombie-situation in w is contained in a situation about which Laura's actual desires are that you shoot her. If we accept that Laura can have desires about non-actual situations, then this seems to derive the correct truth conditions".¹⁹ This will ultimately be very similar

¹⁹As with Blumberg and Holguín, I won't discuss the formal details of von Fintel and Pasternak's account here, though I think they're illuminating. I will make one comment, though. The truth conditions I presented for (22) are really only their first attempt; they revise them subsequently, I think mistakenly. Consider a context for (22) where in fact there will be a cure for zombiehood allowing Laura to live a fine and non-violent life even if she becomes a zombie, although she herself doesn't know this. von Fintel and Pasternak claim that (22) remains true-sounding, but the following sounds false:

(23) If Laura becomes a zombie, she wants you to let her live.

I'm not sure whether (22) is still true in that context; it seems to me to turn on very difficult issues I won't discuss here. But (23) seems to me *definitely* true. After all, it is entailed by the facts, *modus ponens*, and:

(24) If Laura becomes a zombie but there's a cure that would allow her to live a fine and non-violent life, she wants you to let her live.

So either they fail to account for (24), they must deny *modus ponens*, or they must not say that (23) is false. I recommend the last of these courses. But notice that if they go that way, they lose one of the other benefits they claim for their account, that they can make sense of the ill-formedness of counterfactual wishes like:

(25) #If Laura had become a zombie, she wishes you had shot her.

to the account of non-doxastic attitude conditionals that I'll give. Notice there is no first-personal restriction; while this frees them from having to give quite the same sort of explanation that Blumberg and Holguín have to give, they still have to explain the apparent plausibility of any such restriction. As I said, that's one thing I'll try to do.

Another limitation in von Fintel and Pasternak's account, at least as I see it, is that don't wrestle with *why* exactly attitude conditionals in ignorance contexts are problematic, and why we might nevertheless be committed to their truth despite the tensions with knowledge-entailingness and the reasons entailment. *Why* should the attitude conditionals be true in ignorance contexts, and how is that compatible with the intuitions behind these other features of non-doxastic attitude ascription?

And finally, my account is compatible with a position on attitude conditionals with which theirs is not. You may have wondered why I am limiting my attention to *non-doxastic* attitude conditionals in ignorance contexts. Good question; after all, the following also sound true:

- (26) If my application for the position was rejected, I think [/know] that I should have marketed myself differently from how I did.
- (27) If my application for the position was rejected, I'm surprised that they didn't tell me it was rejected at the same time they told everyone else.

Blumberg and Holguín's account on the one hand and von Fintel and Pasternak's on the other would apply equally to doxastic and non-doxastic attitudes. We would then be committed to a *uniform* account of true attitude conditionals in ignorance contexts. But you might think we should treat account for true doxastic and non-doxastic attitude conditionals in ignorance contexts differently from one another, specifically in their entailments. Consider:

- (28) High school was a long time ago, so I'm not sure why he hates me so much, but after talking to people, I've narrowed it down to two possible things: either I bullied him, or I stole his boyfriend. If I bullied him, I regret doing that; and if I

After all, they explain (25)'s infelicity with the same principle with which they would make (23) false.

stole his boyfriend, I regret doing *that*. So, either way, I regret whatever I did to him that made him hate me so much.

That is, frequently, if $\lceil \text{if } \phi, \text{ I } \mathcal{A} \text{ } o_1 \rceil$, $\lceil \text{if } \phi, \text{ I } \mathcal{A} \text{ } o_2 \rceil$, ..., $\lceil \text{if } \phi, \text{ I } \mathcal{A} \text{ } o_n \rceil$ are true in c , and $\lceil \text{I know that, for all } x, \text{ if } x \text{ is } F, \text{ either } x = o_1 \text{ or } x = o_2 \text{ or } \dots \text{ } x = o_n \rceil$ is also true in c , then $\lceil \text{I } \mathcal{A} \text{ whatever is } F \rceil$ is also true in c . That inference, which I'll call an *if-to-ever inference* holds for the non-doxastic attitude expressions I've considered, but not the doxastic ones. Consider, e.g., the following kind of case.²⁰ Call *Homerism* the proposition that there was a unique person who wrote both the *Iliad* and the *Odyssey*, and call *anti-Homerism* the proposition that there was no such person.

(29) I know that either the vast majority of classical archaeologists believe Homerism or the vast majority believe anti-Homerism. If the majority believe Homerism, then I think that there was such a person; but if the majority believe anti-Homerism, then I believe there was no such person. So, I believe whatever the majority of classical archaeologists believe of Homerism and anti-Homerism.

Suppose the speaker has no clue which of Homerism and anti-Homerism it is that the vast majority of classical archaeologists believe. Then the conclusion of (29) is pretty clearly false, I think. Suppose that's right. *Why* should the if-to-ever inference hold for non-doxastic attitude expressions but not for doxastic ones? Something must be different between them. If we continue to retain *modus ponens*, our options seem very limited; it seems we are forced to deny that doxastic attitude conditionals have the simple conditional logical form they seem to have. I'll return to this issue later on. For now, I will note that the view I develop will have the resources to account for the truth of non-doxastic attitude conditionals vs. doxastic attitude conditionals in ignorance contexts; that doesn't seem to be the case for Blumberg and Holguín's and von Fintel's and Pasternak's accounts. That said, I will not account for the difference between them in this paper, since that is too large a task to undertake here. But my view opens room for such an account.²¹

²⁰I've adapted this example from Drucker (2020).

²¹Later on, in section 5, I'll address a case from a reviewer where something of (29)'s general form can be felicitous.

That, then, is the problem of attitude conditionals in ignorance contexts. In section 2, I'll explain why attitude ascriptions with 'ever' free relatives commit us to a face-value account and also solve the problems of unwarranted detachment and absent beliefs. In section 3, I'll explain a lot of the data that remain challenging for face-value accounts like mine. In section 4, I'll distinguish between doxastic and non-doxastic attitude conditionals so as to substantiate the claim that my explanation can allow different treatments of the two, which I claim is a benefit of my explanation. And in the conclusion, I'll describe the work in this area that remains to be done, including doxastic attitude conditionals in ignorance contexts and speech act conditionals.

2 A Face-Value Account

In this section, I'll point to a class of attitude self-ascriptions we *can* make in ignorance contexts, and which entail the attitude conditionals (2)–(6) that I discussed in the previous section. That attitudes or attitude expressions like 'am sad', 'regret', etc. might be knowledge-entailing presents large problems for what I'm calling face-value interpretations of (2)–(6); but I'll argue here that we can show, independent of attitude conditionals themselves, that those principles simply must be false. This argument constitutes important progress to understanding attitude conditionals in ignorance contexts in general.

2.1 Using 'Ever' Free Relatives to Account for Some Attitude Conditionals

My account is motivated by other sorts of utterance that correlate in important ways with attitude conditionals. In an ignorance context, reflecting on mistakes he's made, Todd thinks to himself:

(30) I regret whatever I've done that wasted my time on counterproductive goals.

Suppose this utterance is also true. What truth conditions underlie that? I will argue that they will also make (2)–(6) true as the defender of the face-value account sees them.

'Whatever' as used in (30) is a free relative, specifically an 'ever' free relative (EFR). A lot about the semantics of (E)FRs and the possibly pragmatic phenomena to which

they give rise is controversial. There is nevertheless a constrained-enough space of options to get on with.²²

According to perhaps the most common view, free relatives have the same semantic values as (possibly plural) definites.²³ (EFRs but not FRs, as I'll discuss in a moment, give rise to ignorance and indifference inferences, and perhaps some related inferences.) Thus, 'what I kicked just now' has the same semantic value as 'the thing [that] I kicked just now'. On this view, the following are synonymous:

- (31) a. Chloe owns what I kicked just now.
 b. Chloe owns the thing that I kicked just now.

Assuming a Fregean Heim and Kratzer (1998) *e*-type denotation for definite descriptions and using the iota-operator whose domain is limited to functions of type $\langle e, t \rangle$ that outputs 1 for exactly one input, $\llbracket \text{'what I kicked just now'} \rrbracket^{c,w} = \llbracket \text{'the thing I kicked just now'} \rrbracket^{c,w} = \iota x \text{ thing}'(x)(w) \wedge \text{kicked-just-now}'(x)(\text{me}')(\text{w})$. Then we get $\llbracket (31a) \rrbracket^{c,w} = \llbracket (31b) \rrbracket^{c,w} = 1$ iff $\text{owns}'(\text{Chloe}')(\iota x \text{ thing}'(x)(w)) \wedge \text{kicked-just-now}'(x)(\text{me}')(\text{w})$, i.e., if Chloe owns the thing the speaker kicked just then.

There we assumed a singular denotation, but we can have plural (E)FRs, too; on the view under consideration now, these will have the same semantics as *plural* definites. For example, 'whichever dogs she adopted' will have the same semantic value as 'the dogs [that] she adopted'. Using Link (1983)'s σ operator, we can again deploy a Fregean Heim-and-Kratzer-style *e*-type denotation: $\llbracket \text{'whichever dogs Mara adopted'} \rrbracket^{c,w} = \llbracket \text{'the dogs [that] Mara adopted'} \rrbracket^{c,w} = \sigma x \text{ thing}'(x)(w) \wedge \text{dog}'(x)(w) \wedge \text{adopted}'(x)(\text{Mara})(w)$, i.e., the maximal x such that x is a dog and Mara adopted x , all in w . (σ , like ι , also has a presupposition, namely that there exist a maximal element that satisfies the predicate (here, 'is a dog that Mara adopted') that contains all the other satisfiers as parts.)

The main alternative to this account is to say that they have the semantics of universal quantifiers; a combination of definite and universal interpretations for different kinds of

²²I'm mostly following Šimík (forthcoming)'s very helpful discussion, except where otherwise noted.

²³Pauline Jacobson's work has persuaded a lot of people of this, I think. See especially Jacobson (1995), and also Dayal (1997) and Tredinnick (2005).

(E)FRs is also possible.²⁴ I think on the whole the arguments for the definite view are stronger,²⁵ so I will adopt that analysis going forward. If I switched to the universal view, my arguments would be largely unaffected. Either way, we get a kind of *maximality*: $\lceil I V \text{ whatever is } F \rceil$ will entail $\lceil \text{everything}_i \text{ that's } F \text{ is such that } I V \text{ it}_i \rceil$.

Notice that this is already enough to entail *some* of the attitude conditionals of interest. Given the maximality entailment I just described, (30) seems to entail:

(32) (*As said by Todd in the circumstances*) $\text{Everything}_i \text{ that I've done that wasted my time on counterproductive goals is such that I regret doing it}_i$.

And that's enough to entail the attitude conditional we're after. Let's abbreviate 'my application for the position was rejected' and 'I regret that I wasted the time applying to such a long time' as said by Todd as 'A' and 'B', respectively. First suppose, necessarily, if Todd's application failed and the position was always a counterproductive goal, then his applying was something he did that wasted his time on a counterproductive goal. That's a harmless assumption here. Second, let's suppose Todd's application did fail and his applying was wasting his time on a counterproductive goal. The truth conditions of conditionals are notoriously difficult to specify, but I'll use the heuristic truth conditions for conditionals inherited from Stalnaker and Lewis, according to which $\lceil (if \phi)(\psi) \rceil$ is true iff the closest world in which $\lceil \phi \rceil$ is true is a world in which $\lceil \psi \rceil$ is true.²⁶ So the actual world is one in which Todd's application failed and his applying wasted his time on a counterproductive goal, and, by (32), he actually regrets that he applied to such a longshot. By (strong) centering, the principle that the closest possible world to itself is itself, $\lceil (if A)(B) \rceil$ follows. (Weaker versions of centering aren't relevant here.)

Note that I have assumed there is not a difference between worlds in which $\lceil S \text{ regrets } V\text{-ing} \rceil$ and $\lceil S \text{ regrets that they } V\text{ed} \rceil$. At a minimum, that means I'm assuming $\lceil \text{necessarily, } S \text{ regrets } V\text{-ing iff they regret that they } V\text{ed} \rceil$. Note, e.g., how immediately

²⁴For the universal quantifier view, see, e.g., Bresnan and Grimshaw (1978), Cooper (1983), and Larson (1987).

²⁵This is also Šimík (forthcoming)'s view; see his article for a very good overview of the arguments.

²⁶Here I make what is sometimes called the *limit assumption*, that there exists a unique closest ϕ -world to the actual world. This assumption is questionable, but the way it's questionable wouldn't affect anything I say here.

contradictory the following sounds:

(33) ???Todd regrets getting so drunk, but he doesn't regret that he got so drunk.

The best explanation of why this sounds so bad, I think, is just the necessary equivalence of the two utterances. If that's right, everyone is committed to a face-value understanding of some possible felicitous utterances of (4).

Nevertheless, you might worry that attitude expressions are peculiar in a way that threatens the simple story I'm trying to tell. Frege (1952/1892) drew our attention to apparent substitutivity failures in attitude reports:

- (34) a. Twain's neighbor didn't know that he was Twain.
b. Twain's neighbor didn't know that he was Clemens.

We can easily imagine a situation in which the first of these is true but the second false. In the lingo, 'know' seems to create *opaque contexts*, where you can't swap co-referring expressions in a way that always preserves truth value. It is *incredibly* controversial whether there are any actually opaque contexts beyond quotation.²⁷ If there are, you might worry that (30) doesn't work at all like (32) does, in the way I need it to, since it can seem likely that 'regret' introduces opaque contexts in the way that 'believes' and 'knows' do. The following seem inequivalent:

- (35) a. The literary salon's host regrets that she didn't invite Twain.
b. The literary salon's host regrets that she didn't invite Clemens.

Once again, it's intuitive that we can fill in the details so that the first is true but the second isn't. But if *that's* right, then my reasoning above doesn't work, specifically where I inferred (32) from (30). That inference can look like the following one:

- (36) a. Carolyn believes that every student will do well on the exam.
b. Every student is such that Carolyn believes they will do well on the exam.

²⁷For important work on such matters, see, e.g., Salmon (1986),

It is *extremely* controversial whether this inference works, or when it works if it ever does. It can seem, for example, that (36a) can be true while (b) is false because Carolyn might not *know* of every student that they're a student and thus that they're even taking an exam in the first place.²⁸ Given that, perhaps the reasoning I just went through really does fail, and at the very least I haven't given good enough reason to think it doesn't.²⁹

Seeing how this objection can be overcome will point us toward important facts about attitudes and ascriptions thereof. First, notice that the analogy between belief on the one hand and regret on the other fails at a crucial point. The doxastic analogue of (30) is *not* (36a); rather, the analogue would be more like:

(37) Carolyn believes of whoever is a student [in the relevant class] that they they will do well on the exam.

This, I think, is a stylistic variant of (36b), *modulo* some suggestion of ignorance that I'll return to shortly. But then either we wouldn't infer (35b) from (36), or doing so would be trivial rather than questionable.

I say that (37) is the analogue because (36a) ascribes a doxastic relation to Carolyn, specifically toward one object, the proposition that every student [in the relevant class] will do well in the exam; and what object the complement denotes (or possibly stands for) doesn't vary depending on the facts. (30), by contrast, ascribes a regret relation to an unspecified number of objects; and whatever else is true, it *definitely* doesn't denote or stand for a single proposition. There is nothing like an exportation inference happening in my reasoning above.

Intensional verbs are still strange. We might not have a worry about exportation, specifically, though we should still fear not properly taking into account the intensionality of attitude expressions like 'regret'. But we have to be specific about how intensionality could interfere. First, can we legitimately reason from 'Todd regrets whatever he's done that wasted his time on counterproductive goals' and 'applying to such a longshot [as the

²⁸The literature on "exportation" for 'believe' is large, but see especially Quine (1956), Kaplan (1968), Sosa (1970), Donnellan (1977), Schiffer (1977), Lewis (1979a), Crimmins and Perry (1989), among many others.

²⁹Thanks to a referee for this very important objection.

job he applied to] was one of the things he did that wasted his time on counterproductive goals’ to ‘Todd regrets applying to such a longshot’? It might seem like this is where intensionality might be relevant, but it’s hard to see how. ‘Whatever’ does not allow exceptions, at least when we look outside of embeddings under attitude expressions; if o is an F thing, then $\lceil S V \text{ s whatever is } F \rceil$ is simply untrue so long as S doesn’t V it. Still, maybe attitude expressions introduce important complications that change this.

For the expressions I’ve been focusing on so far, this is just not the case. Since Quine (1956), we’ve known there’s an important difference between the following:

- (38) a. I seek a sloop.
 b. There’s a sloop I seek.

(38a) does not require the existence of any sloops, nor does it require I want any sloop in particular; (38b) does need some sloops to exist, and that I want some in particular. This is also true of verbs of depiction. But note it *isn’t* true of ‘regret’:

- (39) a. I regret a thing I did.
 b. There’s something I did that I regret.

While it may be possible to seek some sloop without seeking any sloop in particular, I don’t think it is at all possible to regret something, but nothing in particular:

- (40) ???I regret something I did last year, but there’s nothing in particular I regret.

Unlike bouletic verbs like ‘seek’ or ‘want’, regret ascriptions with existentially quantified noun phrase complements seem to require there be something in particular the person really does regret. (They might not know what, specifically, it is, of course.) It doesn’t make sense to regret something but not something in particular.

With attitude ascriptions with *universally* quantified noun phrase complements, pretty much every non-doxastic attitude expression seems to work like extensional verbs:

- (41) a. ???I seek every horse in the area, but there might be a horse in the area
 that I don’t seek.

- b. ???I regret everything I said to you yesterday, but there might be something I said to you that I don't regret.

Both of these utterances have the strong feeling of contradiction. But if they really are contradictory, as is the simplest explanation of them (both theoretically and “informally”), then the first conjuncts entail the negations of the second conjuncts, i.e., first that there's no horse that they don't seek and that there's nothing I said that I don't regret. Here's another piece of evidence for the same conclusion:

(42) The artist depicted every horse.

I can think of two ways to make sense of this. First, the ‘every’ is covertly restricted by, e.g., ‘in this room’ or some such predicate. Second, perhaps the speaker means to talk about every *kind* of horse, e.g., the American Paint Horse, the Tennessee Walking Horse, etc. But if the speaker means to use (relatively) unrestricted quantification and isn't talking about kinds of horse but about individuals, (42) sounds basically unintelligible to me because *obviously* untrue.

These facts can escape notice; when philosophers and linguists discuss objectual attitude expressions with quantified noun phrase complements, they tend to focus on existentials. But what I need is not that “exportation” works for existential complements, but rather for universal complements. If there's trouble for my argument, it's not in the step from (30) to (32).

Nor should it be in its other major step so far. That step, recall, is first to go from (32) to ‘I regret wasting so much time by applying to such a longshot’, and to go from there to ‘I regret that I wasted so much time by applying to such a longshot’. If that entailment failed, there should be true utterances of sentences with the following form:

- (43) a. ???*S* regrets *V*-ing, but *S* doesn't regret that they *V*ed.
b. ???Sarah-Lynn regrets getting drunk the previous night, but she doesn't regret that she got drunk last night.

But these kinds of utterance sound crashingly bad. And they're not helped by making the ignorance explicit in a way that would make intensional complications relevant:

(44) ????Though she doesn't know she got drunk the previous night, Sarah-Lynn regrets getting drunk the previous night. But she doesn't regret that she got drunk last night.

Of course, the infelicity of (44) might arise from the first sentence, since now it's stipulated that she doesn't know she got drunk the previous night. But I have already argued that she can regret that in those circumstances, so we can assume at this point that's not what causes the infelicity.

These issues will be with us for the rest of the paper in some guise, but for now I want to complete my argument for the face-value account I favor. To continue, then. That $\ulcorner (if A)(B) \urcorner$ is true in this situation (an ignorance context!) is already fairly striking. I never said anything about Todd's *knowing* that he wasted his time (i.e., that his application failed). The conditional—indeed, its *consequent*—follows simply from (30) and some (I've argued) innocuous stipulations about the case. We're forced to say that regret is not knowledge-entailing in this case regardless of what else we say. This just follows from the ordinary semantics of every FR, including every EFR. It cannot be avoided easily at all once we accept the truth of sentences like (30), at least insofar as (30) doesn't require that Todd know what the things he's done that wasted his time on counterproductive goals actually are.

I've pointed to cases where regret *must* not be knowledge-entailing. Let's return to the problem of absent beliefs. Are there similarly direct violations of the reasons entailment for regret? Consider:

(45) If I my application was rejected, I regret that I applied, because doing so wasted my time.

Suppose the story I've been developing is *explanatory*, in that when some attitude conditionals are true of us in ignorance contexts, they are true of us *because* the corresponding EFR self-ascriptions are true. I have not argued for this, but let's assume it for now. Next, note the explanation of the detached consequent's truth, when detachable, will be that the conditional is true along with the fact that Todd's application was rejected. Next,

suppose Todd’s motivation for the EFR self-ascription (30)—for his general regret—is just that wasting his time is something he does from time to time that he really dislikes doing and feels responsible for doing. Also assume explanation is transitive (i.e., if E_1 explains E_2 , and E_2 explains E_3 , E_1 explains E_3),³⁰ and that ‘because’ is made true because of (sufficient) explanation. Now suppose Todd reasons like follows. Suppose my application was rejected. Then I regret that I applied, and now (because of the considerations just adduced), I can infer that I regret that I applied *because* applying to such a longshot was a waste of my time. And now I can use conditional introduction to discharge my assumption, concluding that *if* my application was rejected, I regret that I applied because doing so wasted my time. Now, I realize this was complicated and depends on some controversial assumptions I haven’t begun to argue for. Still, this makes me optimistic that the challenge posed by the reasons entailment can be met. That said, I want to stress that at this stage, it seems like a harder problem for face-value accounts than knowledge-entailingness, etc. seem to present.

In the next section, I’ll turn to account for *conjunctions* of attitude conditionals in ignorance contexts, which, as we’ll see, are somewhat harder to account for.

2.2 Accounting for Conjunctions of Attitude Conditionals

Using maximality entailments like (32) doesn’t actually get us as far as we might like as far as accounting for the attitude conditionals goes. For suppose Todd thinks:

(46) And if my application *wasn’t* rejected, I regret that I wasted time on having the ideal backup plan.

In this circumstance, he could just as easily conjoin (4) with (46), getting:

(47) If my application for the position really was rejected, I regret that I wasted the time applying to such a longshot, and if my application *wasn’t* rejected, I regret that I wasted time on having the ideal backup plan.

³⁰See Lange (2018) for a recent compelling defense of the transitivity of explanation against “contrastivist” challenges. His version of transitivity is more complicated, but adding the complications wouldn’t change the point.

We can't use the actual truth of *both* antecedents to derive both attitude conditionals, because one of the antecedents must be false. (30) and maximality-entailments like (32) can't account for the felicity of all felicitous attitude conditionals in ignorance contexts.

To get (47), we only need (30) and something like the following:

(48) If my application was rejected, then the thing I did that wasted my time on counterproductive goals was my applying. If my application wasn't rejected, then the thing I did that wasted my time on counterproductive goals was having the ideal backup plan.

(49) If my application was rejected, I regret whatever I did that wasted my time on counterproductive goals, and if it wasn't rejected, I regret whatever I did that wasted my time on counterproductive goals.

We can stipulate (48) as just part of the setup. As for (49), the core idea is this: a person can't *change* whether they regret whatever they did that wasted their time on counterproductive goals based on circumstances external to them that they don't know about. Again, I can simply stipulate that (49) is true. But I don't need to stipulate it, because it follows as the conclusion of an instance of $\lceil \phi$, so $(if \psi)(\phi)$ and $(if \neg\psi)(\phi)\lceil$, which is valid for indicative conditionals. So, so long as it's genuinely true that he regrets whatever he did that wasted his time on counter productive goals, (49) will follow.

And now Todd can reason as follows. "Suppose my application was rejected. Then from (48), the thing I did that wasted my time on counterproductive goals was applying. By (49), I regret whatever I did that wasted my time on counterproductive goals. So, I regret applying; so, I regret that I applied. Discharging the supposition, if my application was rejected, I regret that I applied. Next suppose my application wasn't rejected. Then from (48), my having the ideal backup plan was something that I did that wasted my time on counterproductive goals. But given (49), I regret everything I did that wasted my time on counterproductive goals. So, I regret having the ideal backup plan. So, discharging the starting supposition, if my application wasn't rejected, then I regret having the ideal backup plan. Then using conjunction introduction, (47) follows." Assuming

that conditional introduction is valid for indicative conditionals, this argument will be valid. That, then, is how we obtain the conjunction of attitude conditionals, (47).

Notice I didn't appeal to the specific semantic features of EFRs to derive the conjunction; if each EFR had been replaced with its corresponding definite description, the argument would have worked just as well. I appealed to the fact that for indicative conditionals of the form $\lceil (if\phi)(\psi) \rceil$ such that it's true that ψ and it might be true that ϕ (and thus it might be that ϕ and ψ), it's (of course) true that ϕ —with indicative conditionals, we don't allow actual truths to become untrue unless the antecedent forces us to through incompatibility.³¹ But for the speaker, Todd, it's of course possible both that his application was rejected and that it wasn't (otherwise, he'd never have a motivation for uttering the conditionals even if they could be true). So, given that he actually regrets whatever he did that wasted his time on counterproductive goals, the indicative conditional requires the consequent to be true under a supposition of something possible.

Why, then, is it so much more natural to use EFRs than other definites (assuming, again, that EFRs are really definites)? It's worth looking into that a bit. Compare:

(50) I regret the thing I did that wasted wasted my time on counterproductive goals.

The use of the explicit definite article suggests familiarity, as many linguists have long thought. There are more technical, theory-internal ways of spelling out what it amounts to, but consider Heim (1983)'s classic presentation: “a definite is used to refer to something that is already *familiar* at the current stage of the conversation. An indefinite is used to introduce a *new* referent.” This could be implemented formally through the manipulation of discourse referents for a conversation at a time as registered by the relevant scoreboard.³² While the definite article does differ from EFRs in this way, too, I mean ‘familiar’ in a related but distinct sense. We expect speakers that use definites to have some concrete idea of what answers to the description when there would be no more efficient or briefer way to identify whatever does:

³¹See, e.g., Stalnaker (1975, page 275): “if the conditional is being evaluated at a world in the context set, then the world selected must, if possible, be within the context set as well,” and then a little later, “all worlds within the context set are closer to each other than any worlds.”

³²In addition to Heim, see, e.g., Roberts (2003).

- (51) a. The thing that killed the dinosaurs could happen to us.
 b. Whatever killed the dinosaurs could happen to us.

We expect the first person to be more familiar with what ‘the thing that killed the dinosaurs’ describes than we expect the second person to be. Given that, the EFRs are better in the situations in which we posit robust ignorance. The basic reason is that EFRs have very strong “ignorance inferences”: ‘ $\lceil S \text{ } F\text{s wh-ever is } G \rceil$ ’ often very strongly suggests that S doesn’t know what is F , in some important sense of knowing-‘who’.³³ At a minimum, we can say that many definites don’t suggest lack of knowledge of who answers to that description at all. (51a), for example, has no such suggestion, though (51b) does. But since all the cases we’re focused on will be ignorance cases, and many of them will be cases of quite radical lack of familiarity on the speaker’s part with what falls under ‘ $\lceil \text{wh-ever is } F \rceil$ ’, ascriptions with EFRs as opposed to definite articles send a less noisy signal about whether the person knows what answers to the description or not.

As confirmation of this, stipulate the person *knows* what answers to the relevant description. Imagine that the speaker has some reason to be cagey about what, specifically, they regret of the things they might have done in the pursuit of counterproductive goals. Compare these two utterances, again stipulating he *knows* what satisfies the description:

- (52) a. I regret the thing I did that wasted my time in pursuit of counterproductive goals.
 b. ??I regret whatever I did that wasted my time in pursuit of counterproductive goals.

(b) sounds much stranger than (a). EFRs tend to have either an ignorance inference or what is called an indifference inference: sometimes, when a speaker utters something of the form ‘ $\lceil S \text{ } V\text{s wh-ever is } F \rceil$ ’, they suggest that the object they in fact *V*ed, o as it might happen, was *V*ed just because it was F , and for no further reason that applies to

³³The *locus classicus* of this discussion is Dayal (1997); see also von Stechow (2000), and Condoravdi (2015), among others. While it’s fairly clear ignorance is not an at-issue entailment, there is controversy over whether it’s a presupposition, as von Stechow thinks, or a different sort of consequence of pragmatic reasoning, as Condoravdi thinks.

it, specifically, as opposed to whatever else might easily have been F .³⁴ Since we know the ignorance inference can't hold, and since the ignorance and indifference readings are in near-complementary distribution (at least, where both are *possible*),³⁵ the indifference inference arises. While I think this reading *can* arise, you have to give it quite a bit of background.³⁶ Imagine someone who was constantly drunk during some period of their lives; they've been going back through what they did, trying to find out who they hurt. As they uncover more and more, they come to regret each thing that they did that hurt someone. They think to themselves that they regret starting that big electrical fire, they regret skipping out on a big tab at a restaurant in a way that got their waiter fired, etc. In explaining their process, the person (call them Earl), might say:

(54) I regret whatever I did that hurt someone. When I find it out, I always feel just terrible about having done it. It's the very least I can do, of course.

But without that kind of work in developing a reasonable background, utterances like (52b) sound bizarre. So, all this is to say it's no accident that self-ascriptions in ignorance contexts use EFRs. Definites, at least truth-conditionally, would have worked just as well.

To sum up, I've shown how we can use ordinary EFR attitude self-ascriptions in ignorance contexts to not only derive the truth of corresponding attitude conditionals in context, I've shown how you can derive conjunctions of such conditionals. Finally, I've offered an explanation why these examples sound somewhat better with EFRs than with definites, even though both work fine. So, that's why *every* theorist should accept the truth of face-value interpretations of some attitude conditionals like (4) in ignorance contexts. It's unavoidable.

This raises the question of how general this commitment is. I have shown how (4)

³⁴This way of spelling out the indifference inference is closest to Rawlins (2015), but the counterfactuality descends from at least von Fintel (2000).

³⁵So says Rawlins (2015, page 270), but Condoravdi points to examples of plural EFRs that seem to have neither the ignorance nor the indifference inference. The EFR in (52b) is of course singular.

³⁶If you make it explicit how doing so would satisfy your goals independent of concern you have for particular things you would thereby regret, I think it can work. Thus:

(53) I always try to regret what I do inefficiently, so I can learn not to do that again. So, I regret whatever I did that wasted my time in pursuit of counterproductive goals.

See Rawlins (2015) for more on the importance of teleology to the indifference reading.

and (23) will follow from examples like (30). That leaves open whether we can extend this type of account beyond ‘regret’ to ‘am glad’, ‘am angry’, and ‘hope’, and thereby commit ourselves to a face-value view of the felicity of (2) and (4)–(6). If I succeed, I still need to explain the infelicity of (7) and (8), and of why, e.g., (7) and (8) sound bad even though hoping isn’t knowledge-entailing. I’ll turn to that issue in the next section. For now, it should be pretty clear that the story I was telling for ‘regret’ extends very easily to other attitude expressions. The crucial question is whether ‘am glad’-ascriptions can have EFR-involving complements like with ‘regret’ in (4). They can:

(55) I’m glad about whatever will increase my quality of life significantly.

(56) I’m [angry/very sad] about whatever people have done to mislead me about what I could realistically achieve.

(57) I hope for whatever will just give me peace of mind as quickly as possible.

Focus on (55). This differs from (30): it has the preposition ‘about’ in addition to the attitude expression ‘am glad’. This is, I think, for purely syntactic reasons.³⁷ But now we can go through the same reasoning as above. First, suppose Todd was in fact rejected so that not having to live in Whoville will in fact increase Todd’s quality of life significantly. Then it follows from the maximality of EFRs that he is glad about not having to live in Whoville. Then, by centering, *if* Todd was rejected, he’s glad about not having to live in Whoville. But given someone is glad about not having to live in Whoville iff they’re glad that they don’t have to live in Whoville, the conditional, (3), follows.

In all of these cases, the trick is just to find a relevant EFR-involving attitude ascription that will be true of the given agent. This will depend on the specific circumstances: some agents regret whatever they’ve done that wasted their time on counterproductive goals, of course, but some don’t, if even because they’ve never given the matter any thought. Similar things apply in the case of (56) and (57). To understand why these attitude conditionals are true, we need only understand why some corresponding EFR

³⁷See, e.g., Nebel (2019); for syntactic evidence, he appeals to Dixon (2005). The same point applies to ‘hope for’.

might be true. To be clear, that doesn't establish that the face-value account of (2)–(6) is right. But it does mean that the apparently-strong arguments *against* such an account that I presented in the previous section can't be right; they appeal to principles we must reject, no matter what we say about attitude conditionals. Since the face-value account is, I think, by far the most natural sort of account, this makes the face-value account a kind of “default”. And so far we've seen no reason to move from the default.³⁸ In the next subsection, I'll elaborate on that point a bit further.

2.3 Clarifying the View: Attitude Conditionals Depend on Agent Psychology

A reviewer poses an important challenge to my appeals to EFR data. Answering it will help clarify a really important aspect of the view I'm developing. So, the reviewer suggests the following can sound as good as (30) does:

- (58) It's not true that I regret whatever I've done that wasted my time on counterproductive goals; in fact, I don't regret anything I've done yet because I have no idea which of two things was what I wasted my time on.
- (59) Even if the thing that I did that was wasting time on counterproductive goals was my applying, it's not true that I regret doing it, because for all I know the thing I did that was wasting time on counterproductive goals was actually something else.
- (60) A: When did you begin to regret applying? B: Not until I learned I was rejected.
- (61) I don't regret any of my bets yet, but I know I will regret some of them tomorrow when I see which do and don't end up paying out.

The reviewer suggests that this might pose a problem for my view that attitude conditionals have the simple logical forms they appear to have, since it would seem hard to avoid some mechanism generating the right kind of sensitivity to context.

The view that I've been developing can perfectly well accommodate the fact that in some two circumstances—I *don't* say context, for a reason I'll explain shortly—the same

³⁸Thanks for a reviewer for pressing me on the shape of the dialectic.

type of English EFR sentence can be true and false. The circumstances in which the speaker *really does* have the relevant attitude toward the relevant objects will be the ones in which the corresponding EFR utterance is true. But whether they are true depends on whether the speaker actually has formed the attitude. As an analogy, no one would say the fact that I can say “I believe Biden won the 2020 election legitimately” and someone else can say the negation of that same sentence at all suggests such sentences are sensitive to context beyond in the usual indexical ways (here, pronoun and tense). Whether a person *really has* formed the relevant belief or not is what makes the difference here. Indeed, if epistemic rationality is “permissive”—roughly, given a fixed body of evidence, a rational person could permissibly believe that p or permissibly fail to believe p ³⁹—the same one person could *reasonably* either form the belief or not. Obviously whether the ascription is true depends on whether they’ve done that.

In the reviewer’s examples, Todd seems to be someone who won’t form any kind of “general” regrets, toward whatever (he can’t name the specific things or identify them in different, substantive terms) has certain general properties; he will only regret when he knows, specifically, what he regrets, were he to regret anything at all on the relevant sort of basis. One might deny that this is the right way to go, as far as rationality is concerned. But it is a way many will, I think, find intuitively appealing; and anyway it is surely possible to form one’s regrets with this constraint. People who have this kind of regret-formation habit will make the relevant EFR utterances false of them. And people who have the kind of regret-formation habit I was assuming for most of this section will make the relevant EFR utterances true of them. That means, dialectically, the question of whether I’m right comes in part to whether these facts correspond to the apparently true and false utterances of attitude conditionals.

I think the correspondence turns out exactly right. Consider, e.g., (3) again:

- (3) If my application for the position really was rejected, I’m glad that I won’t have to live in Whoville.

I’ve already argued this will be true in an ignorance context in which (30) is true. But also

³⁹See White (2005).

it seems *false* in a circumstance in which Todd has a habit of only regretting the things he knows he did in some substantive way. That’s why to make the example clearly felicitous the author added “I don’t regret anything I’ve done yet because I have no idea which of two things was what I wasted my time on” to the end of (58). That’s precisely what we would expect on my view. The same is true for “for all we know” and (59) and “not until I learned” and (60). It sounds like this is a person that simply doesn’t form the relevant general kinds of attitudes, in just the way someone might or might not believe that Biden won the 2020 election legitimately. So my view has the resources to accommodate these intuitions, without positing any kind of new context sensitivity. These cases are ordinary dependence on what the corresponding part of the world is like.⁴⁰

That said, you might wonder if an account that *does* appeal to context sensitivity could explain the EFR data better than mine can. Dorr (2011) has argued that belief and knowledge ascriptions involve a lot of sensitivity to context in a way that allows certain otherwise counter-intuitive ascriptions to be true. For example, he argues that if someone knows that the spouse of Hillary Clinton is married to Hillary Clinton if anyone is, and x is the spouse of Hillary Clinton, then, in some contexts it will be true to say that that person knows that x is married to Hillary Clinton if anyone is. This knowledge that x is married to Hillary Clinton if anyone is, Dorr argues, often amounts to a striking and perhaps unexpected piece of *de re a priori* knowledge, surprising because contingent. From that knowledge, perhaps then we can conclude that the person believes *whoever* is married to Hillary Clinton, that person is married to Hillary Clinton. Of course, that must be ‘believes of’ in a certain possibly artificial, philosophical sense of the EFR-belief we’re now attributing; the *de dicto* belief that whoever is married to Hillary Clinton is married to Hillary Clinton is trivial. Importantly, one can explain instances of *de re* belief, knowledge, etc. like this without at all accounting for how *de dicto* attitudes of

⁴⁰I would say the same of the reviewer’s counterfactuals, too:

- (□) Had I not believed that I did some specific thing that was wasting my time on counterproductive goals, then I wouldn’t have regretted anything I did (relevant).
- (◇) ??If my application was rejected but I still didn’t know yet, I still would have regretted applying.
- (□) sounds like what someone who would take themselves to not form regrets in the general way would say of themselves, and (◇) is *not* something such a person would say seemingly truthfully of themselves.

the kind I'm focused on would similarly be possible. Dorr uses true *de dicto* ascriptions to argue for some surprisingly true *de re* attitude ascriptions. But it's surprisingly true *de dicto* ascriptions that I'm trying to account for. So, while I find Dorr's discussion very useful, I don't think it's relevant to my concerns.

For all these reasons, I think everyone is committed to the truth of "face-value" readings of attitude conditionals in ignorance contexts. In the next section, I'll explain away the appearances that they cannot be true in such contexts.

3 *Explaining Away Some Intuitions*

Suppose I'm right: *everyone* is committed to the truth of some attitude conditionals understood as the face-value account would, broadly, have it. That hardly settles matters as far as attitude conditionals are concerned. First, nothing I've said rules out the possibility that some *other* accounts of attitude conditionals (especially in ignorance contexts) might be right for some cases. For all I've said, data might emerge that requires alternative accounts. Still, the work from section 2 does undercut the motivation for these alternative accounts considerably.

Recall these examples:

- (7) ???I have no idea whether my application for the position really was rejected, but I'm very sad that it was.
- (8) ???Todd has no idea whether his application for the position really was rejected, but he's very sad that it was.

To see 'am very sad' patterns with 'regret', recall examples like (56) where 'am very sad' takes EFR complements. The general infelicity of (7) and (8) isn't all that hard to explain. As I said before, one explanation would appeal to KNOWLEDGE NORM OF ASSERTION or one of its cousins. $\lceil S$ regrets that $\phi \rceil$ does not entail $\lceil S$ knows that $\phi \rceil$, but it does entail $\lceil \phi \rceil$ itself. Something stronger seems true: you can't know $\lceil S$ is very sad that $\phi \rceil$ without knowing $\lceil \phi \rceil$ itself. That follows from a version EPISTEMIC CLOSURE: if S knows that p , and p entails q , then S knows q . Then the speaker of

(7) would clearly violate the KNOWLEDGE NORM OF ASSERTION. But this version of EPISTEMIC CLOSURE is implausible, and so it's best to look for a different explanation. It does get something right: self-ascriptions of attitudes can be true but still infelicitous.

We can instead appeal to the fact that 'am [very] sad' (and its very close cousin, 'regret') is a presupposition trigger:⁴¹ normal utterances of $\lceil S \text{ is very sad that } \phi \rceil$ presuppose $\lceil \phi \rceil$. So suggests the projection test:

- (62)
- a. He is very sad that he got drunk last night.
 - b. He isn't very sad that he got drunk last night.
 - c. If he is very sad that he got drunk last night, he's a wise man (but no sage).
 - d. He might be very sad that he got drunk last night.
 - e. She believes that he is very sad that he got drunk last night.

All of (a)–(d) carry the suggestion that he got drunk last night, and (e), doesn't; that's what we'd expect, given 'not', 'if' and 'might' are presupposition holes and 'believes' is a plug. Or consider the 'hey, wait a minute' test:⁴² imagine someone responds to (62a):

- (63) Hey, wait a minute, I had no idea he got drunk last night.

We seem to have a presupposition. But if a speaker says she has no idea whether ϕ , then it won't be common ground that ϕ ; nor would it be accommodated but with extreme difficulty. Thus we should expect $\lceil \text{I regret that } \phi \rceil$ to sound very bad in such circumstances. That's what we find with (7). (Notice 'I have no idea... but if I did..., then I regret that I...' *is* consistent, as we'd expect with presuppositions.) So it's not hard to explain the infelicity of utterances of (7)-like sentences; it doesn't require that expressions like 'am very sad', 'regret', and the rest be knowledge-entailing or even complement-entailing.

The following example will present more difficulties:

- (8) ???Todd has no idea whether his application for the position really was rejected, but he's very sad that it was.

⁴¹This is a long-held view, of course; see, e.g., Kiparsky and Kiparsky (1971).

⁴²See Shanon (1976) for the origins of this test, and von Stechow (2008) for a more recent discussion.

First, notice neither the explanation of (8) that appealed to KNOWLEDGE NORM OF ASSERTION nor the one that appealed to ‘regret’'s being a presupposition trigger help here. Third-personal versions of attitude conditionals sound pretty bad in general. Consider:

(64) ??If Todd’s application for the job was rejected, then he regrets ever applying in the first place.

Call conditionals of the form ‘If ϕ , then S As that ψ ’ *third-personal attitude conditionals*. Imagine (64) said in a context where Todd has no idea whether his application was rejected. Unlike (4), it sounds pretty bad, which is puzzling: why should the first-personal version be felicitously utterable, but the third-personal version (apparently) not?

Infelicities like (8) and (64) might tempt one to think, even if the face-value account must be right for some potential felicitous uses of attitude conditionals in ignorance contexts, perhaps some *other* account will account for (8)’s and (64)’s infelicity. After all, as I said, my account so far is compatible with the possible truth of other accounts, at least applied to different token attitude conditionals (even if it undercuts the motivation for them). There are two problems with this. First, insofar as what I’ve said predicts that there will nevertheless *also* be true readings of (8) and (64), then I still need to explain why (8) and (64) are hard to hear as felicitous; wouldn’t charity suggest we can access the true reading at least somewhat easily?

And second, other accounts that so much as try to capture their infelicity account for it by stipulating a condition that generates these attitude conditionals in ignorance contexts only for the first-person.⁴³ This is *ad hoc*, but also, as I’ll argue a bit later on, there are felicitous third-personal attitude conditionals in ignorance contexts, even for emotive factives. Accounts that write the restriction into the semantics look questionable.

Anyhow, the data here are especially difficult. But for a view like mine, the most important question to ask is whether EFR claims *come apart from* the corresponding third-personal attitude conditionals. So with that in mind, let’s consider some examples. Corresponding to (64), we should check the following:

⁴³See especially Blumberg and Holguín (2019, page 397).

(65) Todd regrets whatever he’s done that wasted his time on counterproductive goals.

Perhaps in this scenario our speaker (Diane again, for simplicity) heard Todd say that to her directly. Now imagine that in the speaker’s next breath, he says:

(66) So, if the thing he did that was wasting time on counterproductive goals was applying to that job, then he regrets doing that [i.e., that he applied to that job].

To my ears, this sounds perfectly fine. The ‘so’, it seems, marks some kind of inference.⁴⁴ We already know at least one inference that might be: it’s one that makes use of the inference. When it’s made explicit that the third-personal attitude conditional is an inference from a corresponding EFR, the utterance of that third-personal attitude conditional is then felicitous. That is the problem I gestured at before for those who say that felicitous attitude conditionals in ignorance contexts must be first-personal.⁴⁵

Even though that is good, why does (64) sound bad when *not* painted as an inference? And can we get felicitous utterances of (8)-like sentences, too?

Before I move on to those questions, I want to stress that (66)’s felicity itself constitutes a pretty strong argument for the view I’ve been developing. If attitude conditionals like (4) are understood in some other way than as entailments of the EFR utterances I’ve been focusing on, then it’s incumbent on proponents of these rival accounts to explain why (66) is felicitous, since no way of understanding it except as an inference occurs to me. It certainly feels like an inference: it would, after all, still be equally acceptable with ‘therefore’ replacing ‘so’. But I won’t press the point here.

Returning, then, to the two questions I just raised, first I’ll address why \lceil if ϕ , then S \mathcal{A} that $\psi \lrcorner$ sounds bad when it’s not marked as the conclusion of an inference. Of course, it’s not the only utterance that tends to sound bad when not marked as an inference. Here are, first, some pretty theory-laden examples:

(67) I might not actually have hands.

⁴⁴For good discussion of inference-marking in natural language, see Pavese (2017, 2019).

⁴⁵Again, I have in mind Blumberg and Holguín. They do say, crediting Harvey Lederman, that there are “echoing” felicitous utterances of third-personal attitude conditionals in ignorance contexts (where someone overhears the speaker assert the first-personal version, and they themselves repeat the third-personal version). Notice inference cases like (66) aren’t like that at all.

In most non-inferential contexts, (67) sounds bad (unless the speaker has reason to worry, or they're in philosophy class). But it *can* be motivated inferentially; consider, e.g.:

(68) An evil demon might be deceiving me into thinking I have a physical body. So I might not actually have hands.

(67) follows from the first sentence of (68) just in case $\lceil \diamond \psi \rceil$ follows from $\lceil \diamond \phi \rceil$ when $\lceil \psi \rceil$ follows from $\lceil \phi \rceil$. Or consider:

(69) The flatscreen's screen isn't actually flat.

(70) The flatscreen's screen isn't anywhere near as flat as my twelve-inch silicon wafer.⁴⁶ So it's not actually flat.

(69) will sound bad if it's a normal flatscreen (not bent or bumpy, e.g.). Or:

(71) The king of France isn't bald.

(72) The king of France isn't bald, because there is no king of France.

Utterances that sound bad outside of inferences can sound good as the conclusions of inferences because inferences can remind us of non-salient *ways* the conclusion can be true or acceptable, ways that might not occur to us otherwise. Thus, people need to be reminded that someone can accept (67) *because* they accept the (epistemic) possibility of a demon deceiver.⁴⁷ And with (69), someone might accept (69) *because* they accept the first sentence of (70), perhaps again because the latter entails the former—whereas otherwise accepting (69) would seem bizarre and unmotivated, almost hallucinatory.⁴⁸ And with (71), it sounds bad when we're not reminded that one way for that utterance to be true or acceptable doesn't require there to be a unique king of France.

For a less theory-laden example, consider:

⁴⁶See <https://news.wisc.edu/curiosities-what-is-the-flattest-thing-in-the-world/>.

⁴⁷Here's one account of what's happening in these cases. In saying (67), we naturally hear the set of worlds quantified over in making the 'might'-claim as including only those worlds compatible with everything the speaker knows. Then we (non-skeptics) tend to treat (68) as bad—false, unreasonable, etc. But when it come as it does in (69), it's clear the 'might' involves quantification over a wider class of worlds, perhaps those compatible with only a subset of what the speaker knows (the knowledge that they would have were they deceived by such a demon, say). For an account of this "intrasubjective flexibility" with epistemic modals, especially in relation to fallibilism in epistemology, see Anderson (2014).

⁴⁸For discussion, see Unger (1975), Lewis (1979b), Kennedy (2007), and most recently, Klecha (2018).

(73) Men can get pregnant.

Even liberal people might see (73) and find it intuitively false or jarring. But now imagine:

(74) Trans men are men who can get pregnant. So, men can get pregnant.

(74) is (typically) fine, in the relevant contexts. Charity should push people to understand (73) as true, but the way it can be true might not occur to people when it's uttered.

There's a *way* for (64) to be true that doesn't occur to people easily when the basic attitude conditional occurs in the third-person. That way is the way I demonstrated in the previous section. If it *did* occur to people more easily, it wouldn't be heard as strange before being reminded of how it can be true. That raises the question of why the first-person attitude conditionals *don't* require our being reminded.

With third-personal conditionals, there's a glaring way for them to be true. Consider:

(75) If she got the job, she's happy about it.

It's hard not to hear that as something like 'if she got the job, she knows and is happy about it'. Why? Conditionals mentioning actions in the consequent *usually* record reactions to the information in the antecedent, very frequently heard as made true by a causal dependency between the events in the antecedent and consequent.⁴⁹ For the process to be causal, the agent must be in some way aware of the event's having taken place.

But 'if ϕ , then I \mathcal{A} that ψ ' cannot be heard that way. After all, 'if ϕ , then I know that ϕ and \mathcal{A} that ψ ' is bizarre; the person shouldn't utter the conditional, then, but just 'I \mathcal{A} that ψ ', or, if they wish to indicate their motivation, 'I \mathcal{A} that ψ because ϕ ': they're more informative given what the speaker knows. So charitable hearers are *forced* automatically to an alternative, non-causal way for 'if ϕ , I \mathcal{A} that ψ ' to be true or acceptable. That, I think, is the origin of the first-personal/third-personal asymmetry, and why it disappears when the third-personal version is the conclusion of an inference. We just more readily think of the alternative way the first-personal version of an attitude conditional can be true or acceptable than we think of the alternative way the third-personal version of an attitude conditional can be true.

⁴⁹See, e.g., van Rooij and Schulz (2019) for this thought and how to make precise sense of it.

Can (8) be felicitous, too? It isn't an attitude conditional; still, the face-value account will entail the truth of some utterances like (8). So, does the infelicity go away, or at least lessen, when (8) is couched as the conclusion of an inference? Let's consider it:

(76) Todd regrets [/is very sad about] whatever he's done that was wasting time on counterproductive goals. He doesn't know this, but his application for the position was rejected. So, he regrets [/is very sad about] applying to that position, even though he has no idea whether his application for the position really was rejected.

This is, I think, not an easy discourse of which to evaluate the felicity. That said, it does not strike me as *infelicitous*. It does strike me as pointlessly drawn-out, so we should imagine a point for saying it all. Thus, imagine someone asks me what Todd's regrets, and specifically whether he regrets applying to the relevant position. Then (76) sounds pretty fine to me. At the very least, whatever oddness it still has for me doesn't seem to go much beyond just how drawn-out the utterance is. But for all relevant purposes this has the structurally relevant features that seemed to make (8) infelicitous.

In fact, (8)-like examples are fairly common with 'want' and are definitely felicitous. To take a recently much-discussed case,⁵⁰ I may be picking up wine for my friends, where the only desirable feature of the wine to me is that my friends like it best. Thus I may deliberate between two bottles, a Malbec and a Pinot Grigio, saying to myself:

(77) If my friends would like the Malbec best, I want to get the Malbec.

I ultimately go with the Malbec. One of my friends, shopping with me but also ignorant as to what the other friends likes, may also say to herself:

(78) If their friends would like the Malbec next, they want to get the Malbec.

Even detaching with *modus ponens* sounds fine with 'want'. Suppose there's an eavesdropper watching me deliberate who knows my friends would like the Pinot Grigio best:

(79) They don't know this, but they really want the Pinot Grigio, not the Malbec.

⁵⁰See Jerzak (2019). See also, well, Plato, especially the *Gorgias* and the *Meno*, and Callard (2017).

On my account, it's not hard to understand why all of these sound fine. The reason these conditionals (or detachments) are acceptable is highly salient, as I've set up the cases. So it's not unexpected that these examples would sound fine.

If we want a unified treatment of these conditionals (I think we should), these examples point to a face-value approach. Otherwise it's tempting to give 'want' a special semantics, for example a relativist one analogous like some treatments of, e.g., epistemic and deontic modals.⁵¹ But as I said, in my view, it's only to be expected that 'want' behaves in this way. After all, it's common ground *why* I want the Pinot Grigio, since it's common ground what my more general desires are. Suppose instead someone says of me:

(80) ???They don't know this, but they really want this piece of garbage here.

An example like (80) will sound bad, since it's not clear *why* I would want the piece of garbage. But suppose instead the person says:

(81) They're making a collage; this is a bottle cap from the drink they need something from. So they don't know this, but they really want this piece of garbage here.

Because the general desire is communicated, (81) sounds fine. We get all this without saying anything interesting or even specific about the semantics of 'want'.⁵²

In sum, none of the data that apparently support thinking that 'regret' and such are knowledge-entailing are inexplicable on a face-value approach to attitude conditionals.

⁵¹See, e.g., MacFarlane (2014) for a very thorough and sensitive statement of relativism, and again Jerzak (2019) for this view of 'want'.

⁵²A reviewer asks why we observe any different behavior for 'want' and, say, 'hope' with respect to the first and third person. I have been arguing that the different behavior is way less than it might have seemed: the third-person ascriptions can be true even in the relevant ignorance contexts as long as we make them the conclusion of an inference that starts with an EFR. But with 'want', we already saw that we rarely need to do that. Is there some general reason why it's easier with 'want'? To be clear I think that's an excellent question, one that might well deserve its own paper. I suspect it's easier to come up with general desires that a person almost certainly has without them telling you explicitly than it is to come up with general hopes that a person almost certainly has. I think that's because 'hope' has a stronger affective component, and thus a weaker association to what a person thinks would be good, in their or the world's interest, etc. For me to have a general hope for something, I have to feel a certain way when I contemplate the possibility. But for me to have a general desire for something, I don't have to feel a certain way, I just have to welcome it, try to make it more likely if I know I can with relatively little cost, etc. So general hopes are harder to ascribe than general desires, since they have this affective component that is very contingent and varies widely from person-to-person. At least, I find that account plausible, though much more needs to be said on the matter. I'm grateful for the reviewer's objection.

Third-personal conditionals and even detachments are sometimes felicitous, and face-value approaches can account for the apparent *inference* involved from moving to EFR-utterances with attitude verbs to the relevant attitude conditionals. Though I think that's an important argument for face-value approaches, this section was largely defensive: data that look bad for that sort of approach don't in fact tell against it.

In the final section, I will fill out one further advantage of my account: because of the kind of explanation I'm giving of non-doxastic attitude conditionals' felicity when asserted in ignorance contexts, my account has room to treat *doxastic* attitude conditionals differently. I'll explain how that's so in the next section, and why it's a good thing.

4 *On Doxastic Attitude Conditionals*

Doxastic attitude expressions can also figure into attitude conditionals uttered felicitously in ignorance contexts. Consider, e.g.:

(82) a. If my application for the position was rejected, I think I'll try to get my application fee back.

b. If my application for the position was rejected, I know I'll try to get my application fee back.

(83) If my application was rejected, I know that my friend's was, too, since he was an even worse applicant than me.

Let's begin by seeing why doxastics are at first glance so tough for we face-value theorists. In short, it's deeply implausible that the detachments with *modus ponens* will be true, even if explainably bad-sounding. If Todd doesn't know his application was rejected, how could he know that he'll try to get his application fee back? Or consider:

(84) If my dog isn't hiding under my bed, then I know no dog is hiding under there.

If their dog isn't hiding under their bed, this could not be felicitously uttered in a context in which the speaker didn't know their dog wasn't hiding under their bed. (By *modus*

ponens, the speaker knows there's no dog hiding there; and then by EPISTEMIC CLOSURE, they know theirs isn't.) But this sentence is most at home in ignorance contexts.

The facts about felicitous use of doxastic attitude conditionals reveal an important linguistic difference: there are examples that must be understood in a non-face-value way, since they are manifestly felicitous in ignorance contexts. As far as what I've argued so far goes, nothing I've said suggests we need a face-value account of doxastic attitude conditionals. That would be a useless observation if the argument that led to face-value accounts of 'regret' and 'am glad' still leads there with 'believes' or 'knows'.

Here, I will first argue that that won't happen: there is no such analogous route to a face-value account of doxastic attitude conditionals. That's because we don't have the EFR-clauses with doxastics that we'd need in order to get the argument going. I'll show that's so for 'believe' and 'know'. I'll also characterize the space of doxastics more precisely, to defend the thesis that doxastic attitude conditionals in ignorance contexts shouldn't be given a face-value account. That will still very much leave open what the proper account of such conditionals actually is. In this paper, however, my aim is only to understand the non-doxastic attitude conditionals. For that project, it will be enough to separate off an account of them from my account of non-doxastic ones.

There are some paradigm cases of doxastics, e.g., 'believe', 'know', 'think', and maybe 'am of the opinion'. Only slightly less obviously, we also have 'am certain', 'doubt', 'am confident', and maybe, at least in "philosophical English", 'have .6 credence', etc. But there are trickier cases, for example 'am surprised', 'am amazed', and 'am stunned'. Since the core of my account is that *true non-doxastic attitude conditionals in ignorance contexts ought (at least sometimes) to be understood with face-value accounts, but doxastic attitude conditionals in ignorance contexts ought not to be*, I must say more exactly what I mean by 'doxastic' and 'non-doxastic'. It's also important because I don't carve this space up precisely the way other theorists do.⁵³

⁵³My distinction is closest to Anand and Hacquard (2013)'s tripartite distinction between informational attitude verbs, preference-based attitude verbs, and hybrids (emotive doxastics). (Their distinction is based on Bolinger (1968)'s distinction between representational and non-representational attitude verbs.) I have a place for some hybrids, like 'am surprised' or 'worry', counting them as doxastic for purposes of the claim I'm defending; Anand and Hacquard's hybrid category is different, though, since it has, e.g., 'fear', whereas mine doesn't have 'fear'.

Doxastic attitudes essentially involve information possession as *part of having the attitude*, whereas non-doxastic ones don't. Here's a sufficient condition: for any sentence ϕ of the form $\ulcorner S \text{ } \mathcal{A} \text{ that } \psi \urcorner$, for any context c , there's a sentence of either the form $\ulcorner S \text{ believes that } \psi \urcorner$, $\ulcorner S \text{ thinks that } \psi \urcorner$, or $\ulcorner S \text{ knows that } \psi \urcorner$ that is informationally equivalent to ϕ relative to c . This also handles:

- (85) a. Diane is certain that Bojack will relapse.
 b. Diane thinks it's certain that Bojack will relapse.

The two sentences in (85) seem equivalent, since affirmations of one simultaneous with denials of the other sound very bad; the information you can take away from one seems to be identical to the information you can take away from the other. Next:

- (86) a. Diane is confident that Bojack will relapse.
 b. Diane thinks it's likely that Bojack will relapse.

Though I think one cannot affirm the (a)-sentence while rejecting the (b)-sentence, one may wonder if what it takes to be confident in a context (assuming it's a normal positive form a gradable adjective, anyway) correlates perfectly with what it takes for a proposition to be likely in a given context. What's important is that being confident that ϕ is gradable *along the same dimension* as thinking it's such-and-such likely that ϕ . For S to become more confident that ϕ is for S to think it's likelier that ϕ ; the following is incoherent:

- (87) #The Dean is now more confident that Jeff will graduate, but he doesn't now think it's likelier that Jeff will graduate.

So it seems we can always mess with the modifiers we apply either to 'confident' or to 'likely' to get an equivalence between (86a) and (86b). (I think the two sentences-in-context are already coordinated, and so this maneuver is unnecessary, but bracket that.)

Anyway, the test also helps with:

- (88) a. Diane has .6 credence that Bojack will relapse.

b. Diane thinks it's .6 likely that Bojack will relapse.

As I said, though, 'credence' (in this use) is probably not ordinary English.⁵⁴

Our condition is sufficient, not necessary, because 'am surprised', 'am amazed', and 'am stunned' wouldn't count as doxastic if it were necessary. There's an affective component 'believes' doesn't have. Should their attitude conditionals get face-value accounts?

To see why not, consider the following example, spoken by Todd:⁵⁵

(89) If my application was already accepted, I'm really surprised that it was.

This can be felicitously uttered in ignorance contexts. But a face-value account will entail Todd is surprised that p even if he has suspended judgment about whether p . That'd be bad. The example even still works if we substitute 'I'm amazed' for 'I'm surprised'.

So, we need one more sufficient condition, and we can allow the necessary and sufficient condition for being doxastic in my sense to be disjunctive between the two sufficient conditions. The second sufficient condition is this: $\lceil S \mathcal{A}$ s that $\psi \rceil$ entails $\lceil S \mathcal{A}$ s to learn that $\psi \rceil$ (relative to a given context c), where $\lceil S \mathcal{A}$ s that $\psi \rceil$ occurs unembedded.⁵⁶ For example, you can't be surprised that p without being surprised to have *learned* that p .

Let's see whether that sorts cases as we might wish so far. Consider:

(90) ???I'm amazed that my application was accepted so quickly, but I'm not amazed to have learned that it was.

Todd is is very hard to understand here. Intuitively, that's because surprise is a reaction not just to facts but facts *as learned*. Here's a different example:

(91) ???I'm amazed that my application was accepted so quickly, but I'm not amazed to have learned that it was.

⁵⁴For this kind of equivalence, see Yalcin (2011).

⁵⁵Blumberg and Holguín (2019) stress the importance of examples with 'am surprised'.

⁵⁶Maybe they must *consider* whether they \mathcal{A} that they themselves learned ψ . I don't think that'd change much. Also, a reviewer helpfully points out that attitude conditionals are counterexamples to this condition without the unembeddedness condition. Of course, there may well be counterexamples even to this condition, so if I'm forced to, I can take the notion as primitive. But I have yet to think of counterexamples to this condition.

Again, being amazed is a reaction to information *as learned*, not just to the information itself. That's why (90) and (91) sound strange. This doesn't reverse: one can be surprised to learn p without being surprised that p , if one doesn't find p surprising but never expected to learn it.

This condition is only useful if it excludes some attitude expressions. It does do that:

(92) I'm afraid that I'm being deceived by all my friends, but I'm not afraid to *learn* that I'm being deceived by all my friends (that'd end the deception).

(92) can be perfectly intelligible if what the speaker most fears of being made a fool of; they don't care that their friends are *trying* to deceive them, but rather that they succeed and the speaker never know it. Fearing that p and fearing learning that p are two different states. And notice we couldn't use a (92)-like example for 'am surprised': a speaker who says 'I'm surprised that ϕ ' must know or believe that ϕ to speak truly.

'I'm afraid that ϕ ' entails or presupposes the epistemic possibility of ' ϕ ' and 'it's not the case that ϕ '. There are other verbs that I would like to say are non-doxastic where we can't appeal to that feature. But we can find other examples:

(93) I'm happy that my friends were giving me a surprise party, but I'm not happy to learn that they were giving me a surprise party (since it ruins it).

(94) I hate that my friends were giving me a surprise party, but I don't hate to learn that they're giving me a surprise party (since it ruins it).

It's easy to give examples like these for many attitude expressions, especially emotives. Surprise is a reaction to information *as learned*, whereas being happy is a reaction to the information itself: you can be happy that you were getting a surprise party; that's happy-making information, even if *as learned* it's bad since it ruins the surprise party.

So call an attitude expression ' \mathcal{A} ' *doxastic* iff at least one of these holds:

- for any $\phi = \ulcorner S \mathcal{A} \urcorner$ that ψ , for any context c , there's a sentence of either the form ' S believes that ψ ', ' S thinks that ψ ', ' S knows that ψ ', or ' S remembers that ψ ' that is equivalent to ϕ given c .

- $\ulcorner S \mathcal{A}$ s that $\psi \urcorner$ entails $\ulcorner S \mathcal{A}$ s to learn that $\psi \urcorner$ (relative to c), where $\ulcorner S \mathcal{A}$ s that $\psi \urcorner$ occurs unembedded.

By the first condition, we get that ‘believes’, ‘thinks’, ‘knows’, ‘is certain’, ‘is confident’, and ‘has credence .6’ are all doxastic; and by the second, we get that ‘am surprised’ and ‘am amazed’ are doxastic. But we also have many expressions that meet neither condition, such as ‘am afraid’ (or ‘fear’), ‘am happy’, ‘hate’, etc. (For these latter expressions, there’s no proposition the belief in or knowledge of which is equivalent to the fearing or hating of.) ‘Am glad’ and ‘am happy’ are relevantly interchangeable, I think, so the same example works for both. You can regret that you made some mistake without regretting to have learned that you did, since you can now fix it. You can be angry that some atrocity happened, say, without being angry to learn that it happened, since you’re happy to learn history. And you can hope that your friends will be giving you a surprise party without hoping to learn that they will. ‘Think’ and ‘know’ are handled by the other conditions. So all the examples come out as they ought to.

Crucially, analogues of (30) with these verbs do not sound felicitous or true in ignorance contexts: they do not generate true- or good-sounding EFR utterances that would commit use to face-value accounts of corresponding attitude conditionals. It is precisely these attitude verbs where we are *not* committed to face-value accounts. That’s good, because face-value accounts with any of these verbs would be unacceptable.

Consider the following examples. Suppose Todd says:

- (95) #I believe [/know/am certain of] whichever of the following claims is the true one:
my application for the position will be rejected, and my application for the
position will be accepted.
- (96) #I am surprised [/amazed] by whichever of the following claims is the true one:
my application for the position will be rejected, and my application for the

position will be accepted.⁵⁷

Examples in both clusters sound awful. They could be uttered truly, though, as (95) would be if Todd really did believe his application for the position was rejected. In that case, we'd need to give him special reason not to say which of the claims he does believe, but that's not too difficult (suppose he has reason to be coy).

A reviewer objects: there are belief ascriptions that seem to be felicitous that have this form. Take 'as far as evolution is concerned, I believe whatever the Bible says'. This might even seem to license doxastic attitude conditionals like 'If the Bible says evolution didn't happen, then I believe that evolution didn't happen'. I think this example expresses an updating plan, to believe p upon learning that the Bible says that p . Thus, where there's *no hope* of finding out the relevant propositions, these utterances sound bad. Imagine there's only one person who witnessed some crime, but they've put themselves in a permanent coma and thus can't communicate. It'd sound absurd to say 'as to who committed the murder, I believe whatever the witness believes'. Or consider 'I believe whichever is true: that dogs were domesticated just once, or that they were domesticated multiple times'; again, in an ignorance context, this sounds awful. The infelicity would make sense if these examples only expressed update plans, because an update plan to defer to the truth would be pointless to express. If it did something else, it would be hard to understand why it sounds bad. If it does express an updating plan, we should expect self-ascriptions of belief that express either impossible or vacuous update policies

⁵⁷A reviewer worries, once we make the cases properly similar, 'am surprised' and similar verbs can take EFR complements as easily as 'regret', etc. If the context entails Todd knows he made some mistake, though he doesn't know what it was specifically, then we can still felicitously say of him 'Todd is surprised by whatever mistake he made'. There's still a large difference here between 'regret' and 'am surprised'. Consider 'I'm surprised by whatever the truth about who built Stonehenge is', said by someone without any candidates in mind, but who knows whoever it is will be surprising, maybe by testimony. This, to me, sounds awful. Someone like that might be surprised by literally nothing, contrary to the surprise self-ascription. So, at a minimum there seems to be a requirement on surprise that the person *at least* have some candidates in mind. But there's no such requirement for 'regret'. The best explanation of this, I think, is that surprise is a response to *learned* information, but here the agent hasn't even learned anything of the form $\lceil X \text{ might have built Stonehenge} \rceil$. Now return to the reviewer's case. When Todd learns he made some mistake, he may have some candidates in mind; in that case, he learned he might have made such-and-such mistake, and he might very well have made such-and-such mistake, etc., and he might be surprised to learn each of those. That that's what he's *really* surprised by (or that he made a mistake at all). So I'm inclined to say while Todd is not surprised by whatever mistake he's made, he is surprised by his making a mistake or that he made a mistake, since those are the things he's learned. That seems to make the best overall sense of the data. I thank the reviewer for their incisive objection.

to sound bad. That's what we find. There are no analogous constraints on 'regret': there are cases where the speaker will never find out what they did, like 'I regret whatever horrible things I did to my friend, and I'm sadder I'll never find out what they were'. I think the original example has a false reading, where it doesn't express an update plan but purports to report current beliefs. But I think it also has a true reading, like 'I believe whatever my accountant thinks about how much I owe the IRS'. This can be paraphrased as 'generally, I believe whatever my accountant thinks about how much I owe the IRS'. This doesn't require in every case I've come to think what my accountant thinks, but if I *never* formed the belief that my accountant expresses to me, then it would be false. But it does admit of exceptions. Similarly, I think the Bible example would be false on every reading if the speaker didn't have *any* beliefs about evolution based on their knowledge of what the Bible says. But one-off cases like Todd's can't be read generically, so there's no symmetry here. In some contexts 'believes whatever' *does* generate the indifference inference I described in section 2, when it's known the speaker will robustly be in a position to execute their update plan, "they always believe whatever their pastor says", e.g. The 'always' seems to me to generate the indifference inference. But we're discussing cases in which the speaker is in no position at all to get such reasons.

The editor also asks a different but related question: suppose Todd tells your friend something, but your friend swears they'll never tell anyone what he said; also suppose we know Todd always tells the truth. The editor wonders whether we could truly say "I believe whatever Todd told you". My intuitions go strongly against this; what we can say, I think, is "I believe whatever Todd told you is true". But these are very different; the latter is just a consequence of my assumption that Todd always speaks truly. The former would involve either a misplaced confidence that I already believe the thing Todd said, though I have no idea what he said and thus would only be making a wild guess, or it would involve the ability to acquire other people's beliefs without having to hear them expressed, which doesn't seem at all possible. It's easy to confuse \lceil I believe whatever is F \rceil and \lceil I believe whatever is F is true \rceil , since generally one believes p if and only if one believes p is true. It's exactly in cases like the one under discussion where one

can believe that some proposition (that one cannot identify in canonical, i.e., sentential terms) is true without believing the proposition itself. So, we have good reason to reject the truth of the ascription, and an error theory about how it can sometimes seem true.⁵⁸

So, I have the means to distinguish doxastic for non-doxastic attitude conditionals, and nothing in my account of non-doxastic ones requires me to extend the explanation to the doxastic conditionals. That's a good thing: only absurdities would result in allowing face-value readings of doxastic attitude conditionals uttered felicitously in ignorant contexts.

5 Conclusion: Future Directions of Research

In this paper, I've tried to account for some very interesting data in a conservative way: once we understand what we're committed to just in virtue of what we need to say about EFR constructions with attitude expressions, we are then committed to the truth of face-value readings of non-doxastic attitudes that are felicitous even in ignorance contexts. This commitment is inescapable, it seems to me, and so we must learn the lessons that such data has to teach us. For example, 'regret' and the rest are not knowledge-entailing.

I would present a formal semantics, particular specifications of the truth-conditions or compositional semantics of the examples of this paper, but I think there's little to learn from doing that in this case: I would introduce no innovations in any formal machinery; rather, I would just use whatever off-the-shelf lexical entries for, e.g., conditionals, EFRs, and attitude expressions that are acceptable for other reasons. The point of my account is that, as far as non-doxastic attitude expressions go at least, no innovation at the level of semantics is needed. The appearance that this kind of innovation is necessary is an illusion created by the relatively indirect way such conditionals become knowably true and thus felicitously utterable. That is, ultimately, the story I tried to tell.

That said there's still a large amount of work left to do. Where I think semantic innovation *is* helpful is in explaining why the doxastic attitude conditionals I had before, (82) and (83), were felicitously utterable in ignorance contexts. I strongly suspect we will want a story that appeals somehow to scope distinctions, but as we saw in section 1 with

⁵⁸I thank the editor and reviewer for their excellent objections here.

example (14), simple wide-scoping strategies are hard to make work:

- (14) I'm very sad that, if my application for the position really was rejected, it was [rejected].

I have some concrete ideas about how such a story would go, but this paper is already quite long, and I think I have made my case adequately. To conclude, I would also like to commend to readers' attention speech act verbs like 'say', 'claim', 'assert', 'order', 'command', 'bet', and 'promise', and 'apologize'. In consequents of conditionals, they all are felicitously utterable in ignorance contexts:

- (97) If I forgot to thank you for driving me to the airport, I apologize for that.
- (98) If I forgot to thank you for driving me to the airport, I promise that I will never forget to do that again.
- (99) If I forgot to thank you for driving me to the airport, I'm telling you I won't again.

These all sound good, but you might worry about the possibility of apologizing for things you don't know you did. Indeed in the case of (99) I think we simply cannot say that its face-value interpretation is can be true in the relevant ignorance contexts. Consider, e.g.:

- (100) If Tom Hanks is going to run as a Democrat in 2024, I'm telling you he will win.

Suppose, bizarrely, Hanks is going to do that. We should not want to say that the speaker told their addressee that Hanks would win. They never asserted that unconditionally.

We thus find a split analogous to that between doxastic and non-doxastic attitude expressions. It seems we should hope to capture this split with the same resources as we used for attitude expressions. This hope is encouraged by the following data pattern:

- (101) I apologize for whatever I did that hurt your feelings so badly.
- (102) ????I'm telling you whichever is true, that Hanks'll run in 2024 or that he won't.

This is extremely suggestive that we're on the right track.

References

- Alvarez, Maria. 2016. “Reasons for Action: Justification, Motivation, Explanation.” In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition. 6
- Anand, Pranav and Hacquard, Valentine. 2013. “Epistemics and Attitudes.” *Semantics & Pragmatics* 6:1–59. 3, 7, 41
- Anderson, Charity. 2014. “Fallibilism and the Flexibility of Epistemic Modals.” *Philosophical Studies* 167:597–606. 36
- Blumberg, Kyle and Holguín, Ben. 2019. “Embedded Attitudes.” *Journal of Semantics* 36:377–406. 2, 8, 34, 43
- Bolinger, Dwight. 1968. “Post-posed Main Phrases: an English Rule for the Romance Subjunctive.” *Canadian Journal of Linguistics* 14:3–30. 41
- Bresnan, Joan W. and Grimshaw, Jane. 1978. “The Syntax of Free Relatives in English.” *Linguistic Inquiry* 9:331–391. 17
- Callard, Agnes. 2017. “Everyone Desires the Good: Socrates’ Protreptic Theory of Desire.” *The Review of Metaphysics* 70:617–644. 38
- Cariani, Fabrizio, Kaufmann, Magdalena, and Kaufmann, Stefan. 2013. “Deliberative Modality under Epistemic Uncertainty.” *Linguistics and Philosophy* 36:225–259. 10
- Condoravdi, Cleo. 2015. “Ignorance, Indifference, and Individuation with *Wh-ever*.” In *Epistemic Indefinites: Exploring Modality Beyond the Verbal Domain*, 213–243. Oxford, UK: Oxford. 26
- Cooper, Robin. 1983. *Quantification and Syntactic Theory*. Dordrecht, Netherlands: Reidel. 17
- Crimmins, Mark and Perry, John. 1989. “The Prince and the Phone Booth: Reporting Puzzling Beliefs.” *Journal of Philosophy* 86:685–711. 19

- Dayal, Veneeta. 1997. "Free Relatives and *Ever*: Identity and Free Choice Readings." In Aaron Lawson (ed.), *Semantics and Linguistic Theory*, volume 7, 99–116. Ithaca, NY: Cornell. 16, 26
- Dietz, Christina H. 2018. "Reasons and Factive Emotions." *Philosophical Studies* 175:1681–1691. 3
- Dixon, Robert M. W. 2005. *A Semantic Approach to English Grammar*. New York: Oxford. Second edition. 28
- Donnellan, Keith. 1977. "The Contingent *A Priori* and Rigid Designators." In Peter French, Theodore E. Uehling, Jr, and Howard K. Wettstein (eds.), *Midwest Studies in Philosophy II: Studies in the Philosophy of Language*, 12–27. Minneapolis, MN: University of Minnesota Press. 19
- Dorr, Cian. 2011. "De Re a Priori Knowledge." *Mind* 120:939–991. 31
- Douven, Igor. 2006. "Assertion, Knowledge, and Rational Credibility." *Philosophical Review* 115:449–485. 4
- Dretske, Fred I. 2005. "The Case Against Closure." In Matthias Steup and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, 13–25. Malden, MA: Blackwell.
- Drucker, Daniel. 2019. "Policy Externalism." *Philosophy and Phenomenological Research* 98:261–285. 7, 10, 11
- . 2020. "The Attitudes We Can Have." *Philosophical Review* 129:591–642. 14
- Frege, Gottlob. 1952/1892. "On Sense and Reference." In Peter Geach and Max Black (eds.), *Translations of the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell. 18
- Gordon, Robert M. 1987. *The Structure of Emotions: Investigations in Cognitive Philosophy*. Cambridge, UK: Cambridge. 3

- Hawthorne, John. 2014. “The Case for Closure.” In Matthias Steup, John Turri, and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, 40–56. Malden, MA: Wiley.
- Heim, Irene. 1983. “File-Change Semantics and the Familiarity Theory of Definiteness.” In Rainer Bäuerle, Christoph Schwarze, and Arnim von Stechow (eds.), *Meaning, Use and Interpretation of Language*, 164–190. Berlin, Germany: de Gruyter. 25
- . 1992. “Presupposition Projection and the Semantics of Attitude Verbs.” *Journal of Semantics* 9:183–221. 5
- Heim, Irene and Kratzer, Angelika. 1998. *Semantics in Generative Grammar*. Malden, MA: Blackwell. 16
- Hyman, John. 1999. “How Knowledge Works.” *The Philosophical Quarterly* 49:433–451. 6
- Iatridou, Sabine. 1991. *Topics in Conditionals*. Ph.D. thesis, Massachusetts Institute of Technology. 11
- Jacobson, Pauline. 1995. “On the Quantificational Force of Free Relatives.” In Emmon Bach, Eloise Jelinek, Angelika Kratzer, and Barbara Partee (eds.), *Quantification in Natural Languages*, 451–486. Dordrecht, Netherlands: Kluwer. 16
- Jerzak, Ethan. 2019. “Two Ways to Want?” *Journal of Philosophy* 116:65–98. 11, 38, 39
- Kaplan, David. 1968. “Quantifying In.” *Synthese* 19:178–214. 19
- Kennedy, Christopher. 2007. “Vagueness and Grammar: the Semantics of Relative and Absolute Gradable Adjectives.” *Linguistics and Philosophy* 30:1–45. 36
- Kenny, Anthony. 2003. *Action, Emotion, and Will*. London, UK: Routledge. 2nd edition, first edition published 1963. 5
- Kiparsky, Paul and Kiparsky, Carol. 1971. “Fact.” In Danny Steinberg and Leon Jacobovits (eds.), *Semantics*, 345–369. Cambridge, UK: Cambridge. 33

- Klecha, Peter. 2018. "On Unidirectionality in Precisification." *Linguistics and Philosophy* 41:87–124. 36
- Kolodny, Niko and MacFarlane, John. 2010. "Ifs and Oughts." *Journal of Philosophy* 107:115–143. 10
- Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1:337–55. 7
- Lackey, Jennifer. 2007. "Norms of Assertion." *Noûs* 41:594–626. 4
- Lange, Marc. 2018. "Transitivity, Self-Explanation, and the Explanatory Circularity Argument Against Humean Accounts of Natural Law." *Synthese* 195:1337–1353. 23
- Larson, Richard. 1987. "'Missing Prepositions' and English Free Relative Clauses." *Linguistic Inquiry* 18:239–266. 17
- Lewis, David. 1979a. "Attitudes *De Dicto* and *De Se*." *Philosophical Review* 88:513–43. 19
- . 1979b. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8:339–59. 36
- Link, Godehard. 1983. "The Logical Analysis of Plurals and Mass Terms: a Lattice-Theoretic Approach." In R. Bäuerle, C. Schwarze, and A. von Stechow (eds.), *Meaning, Use, and the Interpretation of Language*. Berlin: de Gruyter. 16
- Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge. 6
- MacFarlane, John. 2014. *Assessment Sensitivity*. Oxford: Oxford. 11, 39
- Matthews, Gareth B. and Cohen, S. Marc. 1967. "Wants and Lacks." *Journal of Philosophy* 64:455–456. 5
- McGee, Vann. 1985. "A Counterexample to *Modus Ponens*." *Journal of Philosophy* 82:462–471. 9

- Nebel, Jake. 2019. "Hopes, Fears, and Other Grammatical Scarecrows." *Philosophical Review* 128:63–105. 28
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press. 9
- Pavese, Carlotta. 2017. "On the Meaning of 'Therefore'." *Analysis* 77:88–97. doi:10.1093/analys/anx040. 35
- . 2019. "The Semantics and Pragmatics of Argumentation." In Daniel Altshuler (ed.), *Linguistics Meets Philosophy*. Cambridge, UK: Cambridge. 35
- Quine, W. V. O. 1956. "Quantifiers and Propositional Attitudes." *Journal of Philosophy* 53:177–87. 19, 20
- Rawlins, Kyle. 2015. "Indifference and Scalar Inferences in Free Relatives." In Paula Menéndez Benito & Luis Alonso-Ovalle (ed.), *Epistemic Indefinites*, 267–287. Oxford, UK: Oxford. 27
- Regan, Donald. 1980. *Utilitarianism and Cooperation*. Oxford: Oxford. 9
- Roberts, Craige. 2003. "Uniqueness in Definite Noun Phrases." *Linguistics and Philosophy* 26:287–350. 25
- Salmon, Nathan. 1986. *Frege's Puzzle*. Cambridge, MA: MIT. 18
- Schiffer, Stephen. 1977. "Naming and Knowing." *Midwest Studies in Philosophy* 2:28–41. 19
- Shanon, Benny. 1976. "On the Two Kinds of Presuppositions in Natural Language." *Foundations of Language* 14:247–249. 33
- Silk, Alex. 2014a. "Evidence Sensitivity in Weak Necessity Deontic Modals." *Journal of Philosophical Logic* 43:691–723. 10
- . 2014b. "Why 'Ought' Detaches: Or, Why You Ought to Get with My Friends (If You Want to Be My Lover)." *Philosophers' Imprint* 14:1–16. 8

- Sosa, Ernest. 1970. "Propositional Attitudes *De Dicto* and *De Re*." *Journal of Philosophy* 67:883–896. 19
- Stalnaker, Robert. 1975. "Indicative Conditionals." *Philosophia* 5:269–86. 25
- Tredinnick, Victoria. 2005. *On the Semantics of Free Relatives with -ever*. Ph.D. thesis, University of Pennsylvania. 16
- Unger, Peter. 1975. *Ignorance*. Oxford: Oxford University Press. 3, 4, 36
- van Rooij, Robert and Schulz, Katrin. 2019. "Conditionals, Causality, and Causal Propensity." *Journal of Logic, Language, and Computation* 28:55–71. 37
- von Fintel, Kai. 2000. "Whatever." In Brendan Jackson and Tanya Matthews (eds.), *Semantics and Linguistic Theory*, volume 10, 27–39. Ithaca, NY: Cornell. 26, 27
- . 2008. "What Is Presupposition Accommodation, Again?" *Philosophical Perspectives* 22:137–170. 33
- . 2012. "The Best We Can (Expect) to Get?" Paper for a session on Deontic Modals at the Central APA, February 17, 2012. 10
- von Fintel, Kai and Pasternak, Robert. n.d. "Attitudes, Aboutness, and Fake-Restricted Readings." Manuscript. 12
- Šimík, Radek. forthcoming. "Free Relatives." In Daniel Gutzmann, Lisa Matthewson, Cécile Meier, Hotze Rullmann, and Thomas Ede Zimmermann (eds.), *The Companion to Semantics*. Wiley. 16, 17
- White, Roger. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19:445–459. 30
- Williamson, Timothy. 1996. "Knowing and Asserting." *Philosophical Review* 489–523. 4
- . 2000. *Knowledge and its Limits*. Oxford: Oxford. 3
- Yalcin, Seth. 2011. "Nonfactualism about Epistemic Modality." In Brian Weatherson and Andy Egan (eds.), *Epistemic Modality*. Oxford: Oxford. 43