

EPISTEMICS & EMOTIVES

Adam Marushak & James R. Shaw [†]

In the flurry of recent work on the semantics of epistemic modals, it has been noted that they embed under preferential attitude verbs known as emotive doxastics, like *hopes* and *fears* (Hacquard & Wellwood (2012), Anand & Hacquard (2013)). Our goal here is to argue that these embeddings provide an unusually rich source of constraints on the semantics of epistemic modals.

After presenting the basic data, we run through several prominent semantics for modals, focusing first on broadly expressivist positions before transitioning to contextualist forms of descriptivism. We note how pairing each theory we consider with semantics for emotive doxastics yields highly problematic truth-conditions or entailments. The process will help us uncover an increasingly complex data set, with no current theory poised to account for all of it.

Some of this data is important because of how it pressures us to retract influential lessons drawn from modal embeddings under verbs expressing familiar attitudes of acceptance (Hacquard (2006, 2010), Yalcin (2007, 2011)), and related lessons about the nature of *epistemic contradictions* (Yalcin (2007, 2011)). More positively, the data motivate a novel class of views governing epistemic modal domains of quantification that we call *veritic* semantics for modals. On veritic semantics, epistemic modals do not by default function to characterize bodies of knowledge—instead they characterize relevant *truths*. We sketch a contextualist version of this semantics that is poised to account for the total data set, and conclude with recommendations for how expressivists and descriptivists should develop their views if our theory provides the right response to the data.

[†] This paper has benefited from helpful suggestions from Fabrizio Cariani, Andrew Chignell, Jeremy Goodman, Anil Gupta, Valentine Hacquard, Ben Lennertz, Daniel Rothschild, Mandy Simons, Seth Yalcin, audiences at the University of Pittsburgh and Oxford University, and from several anonymous reviewers.

1 THE DATA

Epistemic modals embed under verbs expressing attitudes of acceptance, like *supposes*, *believes*, and *knows*, and facts about these embeddings have played a prominent role in recent debates about the semantics of epistemic modals.¹

What has received less attention, however, is that epistemic modals also embed under so-called emotive doxastics—terms like *hopes*, *fears*, *is concerned*, and *worries*.² Consider the following embeddings (our emphases):

- (1) Many *hope* former agent Dennis Gilbert *might* someday buy the Dodgers.³
- (2) Others *fear* that talking too boldly about American plans *could* fuel a global computer arms race.⁴
- (3) The empty plant still is a hulking presence along the Rock River and some people here *hope there's a chance* it may yet reopen.⁵
- (4) John *fears* that Mary *probably* knew her killer.⁶
- (5) Investors' chronic mistrust of stocks is reigniting *fears* that an entire generation is *unlikely* to stash large chunks of cash in the increasingly unpredictable market as they did in the past.^{7,8}

¹ See Yalcin (2007, 2011), Hacquard (2010) on supposition reports, Hacquard (2006, 2010), Yalcin (2007, 2011), Rothschild (2012), Moss (2013) on belief reports, and Yalcin (2012), Moss (2013) on knowledge reports.

² See Hacquard & Wellwood (2012) pp. 24–25 for a more comprehensive list of emotive doxastics.

³ “Many hope former agent Dennis Gilbert might someday buy the Dodgers,” Kevin Modesti, *Daily News, LA*, Jun. 6, 2011.

⁴ “Cyberwarfare Emerges From Shadows for Public Discussion by U.S. Officials,” Scott Shane, *New York Times*, Sep. 26, 2012.

⁵ “Neighbors like Paul Ryan but may not share his views,” Judy Keen, *USA Today*, Aug. 13, 2012.

⁶ The example is drawn from Anand & Hacquard (2013), who experimentally verify the felicity of translations of this sentence in French, Italian, and Spanish.

⁷ “Invest in stocks? Small players still smarting,” Adam Shell, *USA Today*, May 7, 2012.

⁸ We take it that the modals in these examples all receive epistemic readings. It is widely held that *might* (as opposed to *might have*) is almost invariably epistemic (see Hacking (1967), DeRose (1998), Hacquard (2011), Hacquard & Wellwood (2012), Zeijlstra (2007)). In addition, probability modals are standardly classified as epistemic in the recent literature (see Moss (2013), Rothschild (2012), Swanson (2011), Yalcin (2010)). It is true that probability modals are not classic epistemic modals in the sense of being quantifiers over some set of epistemically accessible worlds. But it would beg the question to thereby exclude them from

Examples like these are not hard to find. An English language corpus-based study in [Hacquard & Wellwood \(2012\)](#) found that proportionately, *might* (a modal that typically receives an epistemic interpretation) embeds under emotive doxastics and emotive factives (e.g. *surprise*) significantly more often than *can* (a modal that typically receives a non-epistemic interpretation and easily embeds in a variety of constructions). [Anand & Hacquard \(2013\)](#) also conducted surveys in French, Italian, and Spanish which indicate that epistemic possibility modals and probabilistic modals like *probable* felicitously embed under emotive doxastics.

In light of this data, any adequate semantics for epistemic modals and emotive doxastics faces the basic task of compositionally generating the correct truth-conditions for these attitude reports and predicting their entailment behavior. Accomplishing this task turns out to be no simple matter. The embeddings in question exhibit complex behavior that departs in important ways from embeddings under familiar attitudes of acceptance. The plan for the next two sections is to run through the leading semantics for epistemic modals, seeing how these frameworks might accommodate the embeddings. We'll argue in §§2–3 that none of the existing forms of expressivist or contextualist semantics can accommodate the data, and we'll use these arguments to motivate a novel form of contextualism in §3.⁹

counting as epistemic modals: [Lassiter \(2011\)](#), for instance, denies that even epistemics like *might* and *must* have a quantificational semantics. One might still question whether the examples contain the uses of epistemic modals that theorists are most interested in capturing. For instance, expressivists might object that the modals in some of our examples are being used to describe objective probabilities, and hence the embeddings in question have no relevance to the project of understanding what [Swanson \(2011\)](#) calls “the language of subjective uncertainty.” We'll argue in §2.1 that expressivists make this move at their own peril: if modals embedded under emotive doxastics receive descriptivist interpretations, then key motivations for expressivism are undermined. One might also object that the modals in our examples are vacuous. In §2.2, we explain why this move is implausible. Finally, while we take the above embeddings to be clearly felicitous, we acknowledge that there may be embeddings of epistemics under emotive doxastics that require a certain degree of contextual setup in order to sound felicitous.

⁹ Our focus will be on expressivism and contextualist forms of descriptivism. We leave out explicit discussion of relativism because we suspect that the distinctive features of the relativist's view are not pertinent to the modal embeddings that are our focus. That is, we expect the relativist's primary options for treating epistemics in attitude embeddings will tend to parallel either the expressivist or contextualist treatments we criticize, in part because the embeddings in question seem to have ordinary, non-relativistic truth-conditions. And, conversely, we suspect the relativist will also be poised to accept relativistic versions of the positive proposals we make for revising expressivist and contextualist theories in §§3–4.

2 EXPRESSIVIST SEMANTICS

Three reasons make expressivist semantics natural starting points for treating the foregoing constructions. First, expressivist views like those of Yalcin (2007, 2011) are tailored to account for peculiarities of epistemics scoped under attitudes. Second, a recent attempt in Anand & Hacquard (2013) to directly account for the data of §1 is formulated in expressivist-compatible terms. Third, our discussion will uncover a close connection between the data from emotive doxastics and the ability of epistemic modals to embed under other epistemic modals, which has received its most thorough treatment in the expressivist framework of Moss (2015).

2.1 DOMAIN SEMANTICS

The semantics we discuss in this paper are formulated in a standard two-dimensional framework, on which expression types are assigned extensions relative to a context and an index. On the domain semantics of Yalcin (2007), the index consists of a world parameter w and an information state parameter s . w ranges over metaphysical possibilities, while s ranges over information states—i.e., sets of metaphysically possible worlds. Epistemic possibility modals like *might* quantify over worlds in s :

$$\llbracket \Diamond \phi \rrbracket^{c,s,w} \text{ is true iff } \exists w' \in s : \llbracket \phi \rrbracket^{c,s,w'} \text{ is true.}$$

That is, *might* ϕ is true at a context-index pair just in case the information state of the index is compatible with ϕ . The information state parameter plays two roles. First, assertions of ‘bare’ modals with unbound information state parameters interact with an expressivist pragmatics of assertion. An assertion of *might* ϕ , in context, determines a property of information states: the property of containing information compatible with ϕ . The assertion of the bare modal functions, roughly, as a proposal for an interlocutor to get into a doxastic state which has this property—i.e. the assertion constitutes a recommendation that one’s beliefs be compatible with ϕ .

Relatedly, and more importantly for our purposes, the information state parameter is bound in attitude reports to ensure that such attitude reports express ordinary truth-conditional propositions. To see how, we can provisionally take a belief-state to be characterized as a set of worlds.

$$\mathcal{B}_A^w = \{w' : w' \text{ is compatible with what } A \text{ believes in } w\}.$$

On the classic Hintikka semantics for belief (Hintikka (1969)), belief reports quantify over this set of worlds. For A to believe that ϕ in w is for ϕ to be true at all of the worlds in \mathcal{B}_A^w . However, if we simply combine this semantics with Yalcin's semantics for modals, A *believes that might* ϕ would be equivalent to *might* ϕ :

$$\begin{aligned} \llbracket A \text{ believes } \Diamond \phi \rrbracket^{c,s,w} \text{ is true} & \text{ iff } \forall w' \in \mathcal{B}_A^w : \llbracket \Diamond \phi \rrbracket^{c,s,w'} \text{ is true} \\ & \text{ iff } \forall w' \in \mathcal{B}_A^w : \exists w'' \in s : \llbracket \phi \rrbracket^{c,s,w''} \text{ is true} \\ & \text{ iff } \exists w'' \in s : \llbracket \phi \rrbracket^{c,s,w''} \text{ is true.} \\ & \text{ iff } \llbracket \Diamond \phi \rrbracket^{c,s,w} \text{ is true} \end{aligned}$$

Yalcin's proposal is to modify Hintikka's semantics so that the attitude report shifts the information state parameter to \mathcal{B}_A^w .

$$\llbracket A \text{ believes } \phi \rrbracket^{c,s,w} \text{ is true iff } \forall w' \in \mathcal{B}_A^w : \llbracket \phi \rrbracket^{c,\mathcal{B}_A^w,w'} \text{ is true.}$$

When no modals are embedded under the attitude, the shifted information state is inert. But when a modal embeds, the ordinary quantificational force of the belief attribution becomes inert, supplanted by that of the modal.

$$\begin{aligned} \llbracket A \text{ believes } \Diamond \phi \rrbracket^{c,s,w} \text{ is true} & \text{ iff } \forall w' \in \mathcal{B}_A^w : \llbracket \Diamond \phi \rrbracket^{c,\mathcal{B}_A^w,w'} \text{ is true} \\ & \text{ iff } \forall w' \in \mathcal{B}_A^w : \exists w'' \in \mathcal{B}_A^w : \llbracket \phi \rrbracket^{c,\mathcal{B}_A^w,w''} \text{ is true} \\ & \text{ iff } \exists w'' \in \mathcal{B}_A^w : \llbracket \phi \rrbracket^{c,\mathcal{B}_A^w,w''} \text{ is true.} \end{aligned}$$

Accordingly, to believe ϕ possible is for ϕ to be compatible with what one believes. Notably, there is no single proposition one needs to believe in order for this to be the case. This feature fits naturally with the expressivist motivations for the semantics, on which no proposition is conveyed by the bare modal construction in assertion.

To accommodate probabilistic modals, an information state can be enriched to consist of something like a probability space. In a simple case, we can take the probability space P to include a probability measure Pr_P defined over sets of worlds. Then we have clauses like the following, where Δ is used for *likely*.

$$\llbracket \Delta \phi \rrbracket^{c,P,w} \text{ is true iff } Pr_P(\{w' : \llbracket \phi \rrbracket^{c,P,w'} \text{ is true}\}) > \frac{1}{2}$$

This semantics essentially says that ϕ *is likely* is true at an information state just in case the state assigns a greater than .5 probability to ϕ . If doxastic states

come graded in terms of degrees of confidence or credences, then they may have something like the structure of a probability measure. To believe ϕ is likely is then to have high credence in ϕ . Again, no ‘modalized truth-conditions’ become the object of the relevant attitude of acceptance here.

In general terms, then, a domain semantics takes modals to “characterize” embedding attitudes in the following way: a modalized construction $\text{MOD}\phi$ determines a property of attitude states, and an attitude report is true just in case the reported attitude has this property.

The question before us is how this semantics for modals fares when combined with semantics for preferential attitudes like emotive doxastics. Importantly, there are two classes of semantics for such attitudes: one which simply extends the classic Hintikka quantificational semantics for attitudes of acceptance just seen, and another which views preferential attitudes as having the distinctive structure of an ordering or ranking. We take them in turn.

On the simplest semantics, a preferential state like a state of hope encapsulates a body of information that can be modeled, at least in simple cases, as a set of possible worlds: those worlds compatible with what the agent hopes. As such, our semantics for *hopes* will parallel that of *believes*.

$\llbracket A \text{ hopes } \phi \rrbracket^{c,s,w}$ is true iff $\forall w' \in \mathcal{H}_A^w : \llbracket \phi \rrbracket^{c,\mathcal{H}_A^w,w'}$ is true, where $\mathcal{H}_A^w = \{w' : w' \text{ is compatible with what } A \text{ hopes in } w\}$.

Accordingly, $A \text{ hopes } \Diamond\phi$ just in case ϕ is compatible with what A hopes.

These are unacceptable truth-conditions. They tell us that if A is completely indifferent to ϕ , then A hopes *might*- ϕ . Suppose that you neither hope Alice comes to your party, nor hope that she doesn’t. You couldn’t care less what she does. If so, then it’s compatible with your hopes that Alice comes. But it in no way seems to follow from this that you hope Alice might come. Indeed, it seems like this is characteristically *not* something you’ll hope.¹⁰

¹⁰ The problem remains even if hoping *might* ϕ additionally requires that ϕ be an answer to a question to which one is sensitive (cf. Yalcin (2011) on believing *might* ϕ). Suppose a friend asks “Who do you hope will show up?” and you reply “I hope Bert comes, but I don’t care whether Alice does.” In this scenario, you’re clearly sensitive to the question of who you hope will show up, and Alice’s attendance is compatible with your answer to this question. But notice *Alice attends* is also a possible answer to the question you are considering. So even if we take into account question-sensitivity, you still count as hoping Alice might come. Some proposals facing similar challenges: According to the dynamic proposal of Veltman (1996), *might* tests the compatibility of an information state with its prejacent. According to that of Willer (2013), *might* functions to transform, or mark, a propositional comple-

The predictions are more problematic for probabilistic modals. If a hope state can, like a doxastic state, be modeled as an information state comprising a set of worlds, then at least we can ask the question of what happens when we characterize that hope state with a possibility modal that expresses a property of information states. By contrast, it is not really clear how a preferential attitude like hope could be characterized by a probability measure in the way that a domain semantics uses probability measures to characterize a doxastic state. Let us explain.

Beliefs, many say, come in degrees. One's levels of confidence in various propositions can be represented by something like a probability measure over them. This idea was essential to getting the semantic values of *likely* or *probable* to engage with our semantics for belief, since on a domain semantics probabilistically modalized complements express a property of probability spaces. Representing the information in a doxastic state as a probability space makes it easy to see how a set of measures could characterize that state. Could some similar story be told for hopes?

Hopes and fears, like all forms of preference, do come in degrees. The problem is that those degrees are irrelevant to the interpretations of modals. The degrees with which one hopes or fears are intuitively a measure of the strength of one's preferences, which varies independently of the likelihood one hopes for or fears. One can fear very intensely that there's still some low probability that the bridge will collapse, or one can hope quite weakly that one is very likely to win some trivial prize. Being told that someone hopes or fears some event is probable does not tell you how strong their hopes or fears are.

If the degrees with which one hopes or fears are not what the scoped probabilistic modals characterize, a domain semanticist seems faced with the task of identifying some other way for probability measures to characterize states of hope or fear, perhaps in identifying some *other* kind of graded structure in

ment as something 'taken seriously' for inquiry in an information state. The former test semantics will generate problems just like a static domain semantics. The latter is challenging to assess since it is not obvious to us how to interpret the informational structure Willer imposes on belief states as a structure that could be applied to a hope state, as the latter is not directly involved in inquiry. But however that proposal goes, *A*'s hoping might ϕ does not seem merely tantamount to ϕ being something *A* 'takes seriously' vis-à-vis the preferential information contained in a hope state. In our most recent case, for example, you may recognize that Alice's coming, or your caring if Alice comes, are important for inquiry or action, without hoping Alice might come. The view will also face the problems connected with nested epistemic modals discussed in §2.3.

those states, while generating plausible truth-conditions for the embeddings. This is clearly a non-trivial task, though there may be many possible avenues to explore. We'll return to a more detailed examination of those avenues in §§2.2–2.3. Our goal for now is just to raise the problems.

Before proceeding, it's worth noting that these issues are not ones the domain semanticist can safely set aside for future research. This is because they apply pressure to a key motivation for domain semantics from the embedding behavior of 'epistemic contradictions' like (6) and (7).

- (6) # It's raining and it might not be raining.
- (7) # It's raining and it's probably not raining.
- (8) # Suppose it's raining and it might not be raining.
- (9) # Suppose it's raining and it's probably not raining.
- (10) Suppose it's raining and it's compatible with what you know that it's not raining.
- (11) Suppose it's raining and, according to what you know, it's probably not raining.

Epistemic contradictions tend to sound infelicitous not only in assertion, but embedded under attitudes of acceptance like supposition. This raises a puzzle, directed primarily at the descriptivist, of explaining why the suppositions sound incoherent while corresponding descriptivist paraphrases can seemingly embed felicitously, as in (10)–(11). The descriptivist arguably must do this without labeling epistemic contradictions *genuine* contradictions, which is in danger of creating a problematic entailment from *might* ϕ or *likely* ϕ to ϕ .¹¹

With the domain semantics above we can explain the embedding behavior as follows: To suppose ϕ is to get into a supposition state with only ϕ -worlds. To suppose *might not* ϕ is to get into a supposition state with some not- ϕ worlds. No coherent supposition state satisfies both conditions simultaneously. And because the bare modal lacks truth-conditions, the danger of positing a problematic entailment can be avoided. A similar story is told for the probabilistic attitudes.

¹¹ See Yalcin (2007) for more detailed discussion.

The problem for this explanation is that embeddings under preferential attitudes are similarly problematic.¹²

(12) # *A* fears that it's raining and might not be raining.

(13) # *A* hopes that it's raining and probably not raining.

What we've just seen is that without some story about how emotive doxastic states can be characterized with a probability measure, the domain semanticist has no account of the infelicity in (13). Worse, even with such an explanation, we cannot transpose the explanation of the infelicity in (8) to that of (12): just as with hopes, fearing might ϕ is not at all to get into an informational fear state compatible with, or otherwise 'indifferent' to, ϕ . Some other explanation is called for.

Note also that any expressivist who leans on the epistemic contradiction data cannot systematically claim that modal embeddings under emotive doxastics tend to involve non-epistemic readings of those modals, or contextual reinterpretations in descriptivist terms, such as interpretations where the modal describes objective probabilities. To do that would be to concede that there is an explanation for the infelicity of (12) and (13) in purely descriptivist terms since, *ex hypothesi*, they must involve descriptivist uses of modals. Whatever that explanation is, it would threaten to generalize to (8) and (9), thereby undermining a key motivation for something like a domain semantics.

This is all just to say: expressivists are not in a good position to deny the embedding data, reinterpret it in descriptivist terms, or to safely relegate it to future research.

Let's return to the main thread of argument. We began by noting that there are two ways to construe the structure of a hope state, and its resulting semantics. The first, which takes a hope state to be an ordinary information state,

¹² It is an open question what consistency requirements hopes may be subject to. It may also be worth distinguishing between, on the one hand, the question of whether hoping ϕ and hoping $\neg\phi$ exhibits some kind of incoherence, and, on the other, whether hoping $\phi \wedge \neg\phi$ does. The latter, for example, seems much more clearly problematic. We don't want to presuppose answers to any of these questions at this stage. The fact is that whatever one says about these issues, (12)–(13) are marked. We're claiming that it would be extremely surprising if the odd-sounding character here wasn't connected to the odd-sounding character of epistemic contradictions scoped under attitudes of acceptance (some of which are arguably also not subject to clear standards of consistency, like supposition). We're grateful to [blinded for review] for pressing us to clarify this issue.

faces the obstacles outlined above. But there is another construal of preferential states which takes them to be fundamentally different in structure from ordinary representational states, or attitudes of acceptance, like belief and supposition. According to this rival view, preferential states like desires are inherently structured as *preference orderings* (Bolinger (1968), Stalnaker (1984), Heim (1992), Villalta (2000, 2008)). Instead of having the structure of a set of propositions or worlds, a hope state would be structured as a ranking of propositions. To hope ϕ , on this view, could involve being in a preferential state situating ϕ high enough in a ranking relative to a contextually determined set of alternative propositions. Such a semantics might look roughly as follows, where S_c is a contextually determined set of alternative propositions, $\llbracket \phi \rrbracket^c = \{w : \llbracket \phi \rrbracket^{c,w} \text{ is true}\}$, and $>_{\mathcal{H}(A,w)}$ is a preferential ranking on propositions supplied by the preferential state connected with A 's hopes in w .

$$\llbracket A \text{ hopes } \phi \rrbracket^{c,w} \text{ is true iff } \forall q \in (S_c - \llbracket \phi \rrbracket^c) : \llbracket \phi \rrbracket^c >_{\mathcal{H}(A,w)} q.$$

To hope ϕ , on this view, would involve preferring ϕ to all of the contextually relevant alternatives. However, Anand & Hacquard (2013) have noted a key obstacle for such a view in accounting for the data of §1. The problem is essentially a more general version of the structural problem from probabilistic modals noted above. On the domain semanticist's view, modals function in attitude reports to characterize the informational structure of the embedding attitude state. The problem, if modals inherently characterize states of information, is that there is a mismatch between the characterizing function of the modal, and the structure of the state it is supposed to characterize. For example, on the current proposal an embedded *might* serves to quantify over worlds in an information state supplied by the embedding preferential state. But if the preferential state does not supply a set of worlds to quantify over, rote application of the compositional machinery will effectively crash, either generating trivial truth-conditions or no truth-conditions at all.¹³

¹³ For example, extending the foregoing semantics for *hope* to a domain semantics, along the lines we took to extend the semantics for *believes*, would involve adding an information state parameter whose value is shifted by the attitude verb to the 'information' carried in the attitude state—now seemingly given by the preferential ranking $>_{\mathcal{H}(A,w)}$. We could then seek to rank the set of worlds given by $\llbracket \phi \rrbracket^{c,s} = \{w : \llbracket \phi \rrbracket^{c,s,w} \text{ is true}\}$ in the hope ranking. In the case of *A hopes might* ϕ this would yield the following.

$$\begin{aligned} \llbracket A \text{ hopes } \Diamond \phi \rrbracket^{c,s,w} \text{ is true iff} \\ \forall q \in (S_c - \llbracket \Diamond \phi \rrbracket^{c,s,w}) : \llbracket \Diamond \phi \rrbracket^{c,s,w} >_{\mathcal{H}(A,w)} q. \end{aligned}$$

It is worth stressing that this is not merely a compositional problem, but one connected with metaphysical questions about a hope state's structure, and the truth-conditions of embeddings. If *might* ϕ expresses compatibility with an information state, and *likely* ϕ characterizes a probabilistically structured state, what function can they serve in characterizing a ranking of propositions? What truth-conditions would plausibly result from applying these constructions to a preference ordering, even if they could be forced to compositionally engage?

This is just to say that the issues we've uncovered for a domain semantics are not helped, but exacerbated by adopting a conception of a hope state as comprising a preference ordering on propositions. There is non-trivial, and very important work to be done in explaining the embeddings.

So far, this is just to set up the problem. Where to go from here? There are several avenues to pursue. First, we could take on the task of giving an entirely new kind of structure to a hope state, so that it is neither comprised by a set of worlds, nor by a preference ordering on propositions. Second, we could change the characterizing function of the epistemic modals or the 'content' that gets assigned to the modalized clauses. (And of course, we could do both together.) Third, we might try to eliminate the worries by claiming that modals are vacuous in these reports.

We'll look at all these proposals eventually. But it will be most helpful to begin with a fourth inventive option embraced by [Anand & Hacquard \(2013\)](#). This is to concede that there is a fundamental mismatch between the characterizing function of modals and preferential attitude states, insofar as emotive doxastic states comprise preference rankings, but to deny that emotive doxastic

This effectively leads to the aforementioned compositional crash. $\Diamond\phi$ tests ϕ for compatibility with an information state like a set of worlds. But $>_{\mathcal{H}(A,w)}$ is a ranking, not a set of worlds. As such $\llbracket \Diamond\phi \rrbracket^{c, >_{\mathcal{H}(A,w)}} = \emptyset$, irrespective of the value of ϕ . Even if this didn't lead to infelicity via a presuppositional constraint of non-triviality on epistemic modal bases ([Geurts \(2005\)](#)), the resulting truth-conditions for hope reports would clearly be unacceptable. Anand & Hacquard formalize the problem here for purely preferential attitudes by conventionally setting the shifted information state to \emptyset to reflect that there is no 'true' information state for a purely preferential state to supply (bearing in mind the caveat that, as we'll soon discuss, Anand & Hacquard do not treat *hopes* as a purely preferential attitude). This is claimed to systematically result in the violation of a presuppositional constraint of non-triviality for epistemic modal bases. They note that one could perhaps also generate related problems if no shift occurred at all, since $\llbracket \Diamond\phi \rrbracket^{c,s}$ is independently guaranteed to be a tautology or contradiction, regardless of the values of c or s , owing to the insensitivity of $\Diamond\phi$ to the world parameter.

states are *merely* preferential states.

2.2 HYBRID-STATE SEMANTICS AND MODAL VACUITY

Emotive doxastics are so-called because their preferential component seems bound up with separate doxastic attitudes. Linguists have noted that hoping ϕ or fearing ϕ seems to entail that one believes both ϕ and $\neg\phi$ are epistemically possible (Portner (1997), Truckenbrodt (1997), Scheffler (2008), Falaus (2010)). Philosophers working on hope have come to the similar conclusion that hoping ϕ requires active agnosticism about ϕ (Bovens (1999), Pettit (2004)). Perhaps what this tells us is that hopes and fears are really *hybrid* states, comprising both a preferential and a doxastic component.

Anand & Hacquard note that this hypothesis can be used to explain how epistemic modals embed under emotive doxastics. The idea is that in a report *A hopes might ϕ* , the modalized complement *might ϕ* is used to characterize only the doxastic component of the hybrid state, while only the unmodalized preja-cent ϕ characterizes its preferential component. More generally, for epistemic modals MOD, hoping MOD ϕ involves believing MOD ϕ and hoping ϕ . Since the doxastic state is assumed to have the familiar structure of an attitude of acceptance, there is no obstacle to having it characterized by the modal complement. And since the modal is inert when characterizing the preferential state, no issues of compatibility arise in characterizing the preferential component either.

The resulting semantics (whose compositional details aren't needed for our discussion) not only resolves the key puzzle of the previous section in this way, but has three further important virtues.

First, it overcomes the problem of explaining the asymmetry between belief and hope noted in §2.1. Whereas believing *might ϕ* is connected to agnosticism about ϕ , hoping *might ϕ* constitutively involves hoping ϕ , which is why hoping *might ϕ* is not characteristically bound up with preferential indifference.

Second, the view opens up an explanation of the unembeddability of epistemic contradictions: hoping ϕ and *might not ϕ* involves hoping both ϕ and hoping not ϕ . Arguably, reporting contradictory preferences could be to report an aberrant, confused state of mind.¹⁴

Finally, Anand & Hacquard's theory explains one last bit of puzzling data:

¹⁴ Though note that this account requires something like the incoherence of both hoping ϕ and hoping $\neg\phi$ (as opposed to the clearer incoherence of hoping $\phi \wedge \neg\phi$). See n.12, n.37.

while possibility modals like *might* embed neatly under emotive doxastics, necessity modals like *must* seem to resist embedding. It is not easy to get a reading of (14) where *must* functions as an epistemic modal.

(14) ? I hope that Alice must be coming to my party.

The resistance is borne out by corpus studies and surveys in Romance languages.¹⁵ The semantics accounts for this as follows: Hoping *must* ϕ involves (through its preferential component) hoping ϕ , which in turn (qua emotive doxastic) requires believing might not- ϕ . But hoping *must* ϕ also involves (through its doxastic component) believing *must* ϕ . But believing *must* ϕ is incompatible with believing might not- ϕ . Hence the infelicity of the original attribution.

Despite these many virtues, we believe the proposed hybrid semantics has a simple, and fatal, problem.¹⁶ Constitutively, the view licenses two entailments for epistemic modals MOD.

$A \text{ hopes MOD } \phi \models A \text{ believes MOD } \phi$

$A \text{ hopes MOD } \phi \models A \text{ hopes } \phi$

But both entailments fail.

Consider (15a) said by a dejected boyfriend to a confidant shortly after being dumped.

- (15) (a) I hope we might get back together.
(b) I believe we might get back together.

Or consider (16a), said by someone after overhearing that a friend is trying to persuade some celebrities (whose identities are as yet unknown) to come to their surprise party.

¹⁵ Again, see [Hacquard & Wellwood \(2012\)](#), [Anand & Hacquard \(2013\)](#).

¹⁶ Anand & Hacquard only offer a semantics for embeddings of modal auxiliaries like *might* and *must*, so strictly speaking, some of the examples to follow actually challenge an extension of their semantics to cover all epistemic modals. However, there are significant obstacles to treating other modals on a different model. Anand & Hacquard's own experimental studies indicate that probability modals felicitously embed under emotive doxastics. And we give examples below illustrating the felicity of embeddings involving modal adjectives. All of these cases present the same sort of compositional mismatch problem that motivates Anand & Hacquard's original proposal. So there is pressure for them to treat these other cases the same way, having the modal characterize only the doxastic state connected with the emotive doxastic attitude, and thereby predicting that *hopes MOD* ϕ entails both *believes MOD* ϕ and *hopes* ϕ .

- (16) (a) I hope Brad Pitt might be coming to my surprise party.
 (b) I believe Brad Pitt might be coming to my surprise party.

Following up (15a) with (15b), or (16a) with (16b), seems to provide news that the speaker takes their hopes to be fulfilled. Note also that saying (15a)[(16a)] *after* saying (15b)[(16b)] sounds odd, as if one is contradicting or taking one's earlier assertion back. In other words, the entailments from (a) to (b) in both cases are blocked.¹⁷

The entailment for probabilistic modals is just as problematic. Suppose a group of scientists is sending off a fly specimen to be examined by an expert. Compare:

- (17) (a) We hope the specimen is more likely to belong to G. hackmani than to G. balachowskyi.
 (b) We believe the specimen is more likely to belong to G. hackmani than to G. balachowskyi.

(17a) can surely be true without (17b): suppose the scientists would like to have a member of G. hackmani in their collection but have no idea about the genus of their specimen.¹⁸

The entailment also fails for modal adjectives embedded under quantifiers.

¹⁷ In saying that the entailment is blocked, we mean that the proposition expressed by the first sentence does not entail the proposition expressed by the second, relative to the same contextual contributions to the interpretation of the modal. We appear to have constant contextual contributions of this kind, precisely when the sentences are taken in immediate succession. One might wonder whether Anand and Hacquard could consistently explain these blocked entailments by holding that the doxastic component of hope is weaker than the state of mind reported by overt talk of belief. But this would leave us unable to explain the very data that motivate positing a doxastic component of hope in the first place. For example, sentences with the following forms are marked:

- (i) # Tanya hopes that Samir is in his office and she doesn't think that he might be there.

This provides evidence that the doxastic attitude borne to p in virtue of hoping p should at least entail believing p epistemically possible. Note that this example also casts doubt on the idea that the connection between hope and agnosticism is merely pragmatic—for example being generated via an implicature. If it were, one would expect the above sentence to simply cancel any such implicature. We're grateful to [blinded for review] for suggesting these ideas as responses to consider.

¹⁸ The modalized complement in these sentences comes from a dialogue discussed in Moss (2013). Note that, as Moss presupposes, this doesn't seem to describe an objective probability: the specimen either belongs to a given genus or it doesn't, so the objective probability

(18) We hope that all of the applicants are possible hires.

(19) We believe that all of the applicants are possible hires.

One is warranted in asserting (18) without even looking at any of the applicants' files. Not so with (19). And again, following (18) with (19) seems to provide news that the speaker's hopes were fulfilled.¹⁹

Indeed, the data reveals a noteworthy pattern. Not only does it seem like there is a failed entailment from *hopes MOD* ϕ to *believes MOD* ϕ but, if anything, there seems to be an *anti*-entailment. To say one hopes that something might be the case, or is likely to occur, seems precisely to express doubt in the complements—the conscious withholding of belief or other doxastic commitment. In a way, this is to be expected. We appear to have on our hands a manifestation of a general phenomenon already noted: that hope and fear require agnosticism. The data suggests that this agnosticism extends to modalized complements. As such, Anand & Hacquard's proposal is actually getting the characterization of the doxastic attitudes backwards.²⁰

The second constitutive entailment of Anand & Hacquard's view—from *hopes MOD* ϕ to *hopes* ϕ —is in slightly better standing. As already noted, there seems to be a tight connection between hoping-possible and hoping-likely on the one hand, and simple hope on the other. The problem is that the connection seems defeasible. I can hope your keys might be in your new car because if the keys might be in the car, I'll get to see the inside of it. As such, I might be indifferent as to whether the keys are actually in the car: I'm most interested in tagging along with you to look for them in certain locations. Similarly a club promoter can hope that a celebrity might be showing up because the

ities in question would have to be either 1 or 0. Could the probabilities be frequencies or propensities instead? The problem with the former is that the probabilities in question are single-case. The problem with the latter is that it seems confused to speak of the propensity of a specimen to belong to a genus. Could the scientists simply be hoping that the expert believes the specimen is more likely to belong to G. hackmani? We doubt it: the scientists are interested in this belief only insofar as it is true, since their ultimate goal is to obtain a member of G. hackmani. And, again, even if any of these descriptivist reinterpretations were in force, it would raise the problems for expressivists discussed on p.9.

¹⁹ The complements in these embeddings are based on examples from Swanson (2011) in which a quantifier takes wide scope over an epistemic modal. See also Moss (2015) for discussion of similar examples involving modal adjectives.

²⁰ One important improvement to the view would be to take modals to target an implicit doxastic state not of belief, but of believing-possible. To be an improvement, though, this would importantly require nested epistemic modals to be non-redundant—otherwise we get the same problematic entailments again. We discuss the issue in §2.3.

mere possibility of his attendance brings in crowds. But the promoter might be indifferent as to whether he actually appears.²¹

Again the problems are just as bad for probabilistic modals. I can hope each of two teams in a sports match have equal chances of winning (and losing) because I'm more interested in a good match than the outcome. This doesn't require me to simultaneously hope of each team that they both win (and lose).²² A student can hope a snowstorm tomorrow is likely so that school is proactively closed. Whether it actually snows might not matter to her—she may just want the day off since she hasn't finished her homework. Indeed, she might even prefer a sunny day off more than a snowy one.

As before, the problem also arises for modal adjectives embedded under quantifiers. If modals play no role in characterizing the preferential component of an emotive doxastic, then *A hopes that* $(\forall x)(\Diamond Fx)$ should entail *A hopes that* $(\forall x)(Fx)$. But this entailment clearly fails: you can hope that all of the job applicants for a single position are possible hires without hoping (confusedly) that all of the applicants get hired.

The entailments that we've claimed fail are not easily separable features of Anand & Hacquard's view. Indeed, they seem to constitute the very heart of the proposal, being precisely what secures the four virtuous features of their view discussed above. If the hybrid semantics can be resuscitated, it seems like it will require important changes that are in danger of upsetting the virtues of the framework along with its vices.

Before moving on, it is worth considering an idea floated earlier that we should now see is increasingly unappealing: the view that epistemic modals are vacuous when embedded under emotive doxastics. The view may have seemed credible for possibility modals. Hoping Alice might come to the party seems scarcely distinguishable from hoping Alice does come. But there are two reasons to reject this proposal.

²¹ Could the promoter merely be hoping that people believe the celebrity might show up? Not necessarily: suppose the promoter hopes the celebrity might show up because people would be angry if it later turned out that there was no chance the celebrity would appear.

²² One might object that *chance* is typically used to describe objective probabilities, so our examples do not contain the uses of modals expressivists aim to capture. But we think treating this as a default reading of *chance* is more of a philosophers' convention than a feature of natural language. E.g., Swanson (2011) includes talk of chance in his examples of the language of subjective uncertainty. Moreover, even were this subject to a descriptivist reinterpretation in terms of objective probabilities, it would raise the problems for expressivists discussed on p.9.

First, if epistemic modals are always vacuous in these embeddings, then $A \text{ hopes MOD } \phi$ should be equivalent to $A \text{ hopes } \phi$ —i.e. the former should entail the latter, and vice versa. Our examples above already demonstrate the failure of the predicted left-right entailment. There are grounds for doubting the right-left entailment as well. Anand & Hacquard note that consultants in informal discussions reported that *might* as used in (20b) indicates that John takes the possibility that someone is listening in to be more remote.

- (20) (a) John is worried that someone is listening in.
 (b) John is worried that someone might be listening in.

If *might* is completely semantically inert (as even Anand & Hacquard do not claim) it will become much more difficult to account for these intuitions.

Second, there is evidence against even the weaker thesis that epistemic modals are *sometimes* used vacuously when embedded under emotive doxastics. We noted above that epistemic necessity modals resist embedding under emotive doxastics. But if epistemic modals can sometimes embed vacuously, one wonders why we cannot access these vacuous readings to save $A \text{ hopes must } \phi$ from infelicity.²³ The only remaining option is hold that *might* and *probably* alone are capable of receiving the proposed vacuous readings. But this proposal simply appears *ad hoc*.

All should agree that $A \text{ hopes/fears might } \phi$ has strong connections to $A \text{ hopes/fears } \phi$. But the data reveals that the best explanation of these connections does not treat even possibility modals as semantically inert. Indeed, such modals arguably don't even sound vacuous. Rather, as Anand & Hacquard's consultants reveal, they seem like a way of hedging the doxastic commitments of an attitude holder. What needs explanation is what such hedging comes to, and why it is compatible with, and may often come along with, corresponding non-modalized hopes and fears. To give such an explanation, though, we need to work on the assumption that possibility modals make a non-trivial semantic contribution in the emotive doxastic constructions that embed them. Let's continue to explore what that contribution might be.

²³ In fact, the only cases in which these embeddings sound reasonably felicitous are those in which the modal clearly makes a semantic contribution: e.g. one might hope that someone *has* to be the murderer so that there is no doubt about his guilt. See §3.2 for further discussion.

2.3 NESTED MODALS AND PROBABILISTIC CONTENT

We've just seen that $A \text{ hopes } \text{MOD}\phi$ does not entail $A \text{ believes } \text{MOD}\phi$. And we noted that this is just what we would expect given that $A \text{ hopes } \phi$ seems to entail a kind of agnosticism about ϕ . Indeed, as linguists and philosophers have all noted, $A \text{ hopes } \phi$ seems to involve believing both ϕ and its negation are epistemically possible. Why not extend this treatment to the modalized complements, as the data seems to suggest?

There is an obvious obstacle for this approach. If $A \text{ hopes } \phi$ entails $A \text{ believes } \Diamond\phi$, and nested modals are equivalent to the inner modal, then we'll have the problematic entailments of the previous section anyway: $A \text{ hopes } \Diamond\phi$ will entail $A \text{ believes } \Diamond\Diamond\phi$, which will entail $A \text{ believes } \Diamond\phi$. The treatment of nested modals as redundant in this way is commonplace in expressivist semantics. In Yalcin's semantics, for instance, the outer modal introduces redundant quantification over the information state parameter, while in dynamic proposals like Veltman (1996) and Willer (2013), updating on a nested modal is equivalent to updating on the inner modal.

Moss (2015) objects to this treatment of nested modals, arguing that it fails to capture the semantic differences between sentences like (21)–(23):

- (21) Bob is a possible hire.
- (22) Bob is probably a possible hire.
- (23) Bob is definitely a possible hire.

She goes on to develop an expressivist semantics in which nesting epistemic modals can indeed have important semantic effects. As a result, her semantics seem well placed to avoid the entailment from $\text{hopes } \text{MOD}\phi$ to $\text{believes } \text{MOD}\phi$ by denying the redundancy of nested modals, which would offer an improvement over the accounts we've considered so far.

The central idea of Moss's semantics is that unembedded epistemic modals express constraints on credences. We've already seen something like this proposal in Yalcin's semantic entry for *likely* (see also Rothschild (2012)), but Moss departs from other expressivists in taking the semantic value of an epistemically modalized sentence to be context-sensitive. On her semantics, context supplies a partition of logical space, and epistemic modals place constraints on credences, conditional on elements of the contextually supplied partition.

For example, consider *Bob might be the best candidate*. Here the context of use might partition the logical space as follows: we can separate the worlds in which experience is the most important virtue in a candidate from the worlds in which it isn't (each set of worlds constitutes a proposition). To accept the sentence is to accept that there is some proposition in this partition that would lead you to accept that Bob is the best candidate. More formally, the semantic value of the sentence is the set of probability measures for which there is some proposition in the partition conditional upon which the measure assigns probability 1 to Bob's being the best candidate.

As before, many of the details of this semantics won't be necessary to appreciate how it engages with emotive doxastic embeddings. There are two significant features of Moss's semantics for our purposes. First, she allows context to supply different partitions for each of a pair of nested modals, in which case embedding modals under other modals can have a non-redundant semantic effect.²⁴ Second, the semantic values of modalized claims are now treated not as sets of worlds, or measures, but sets of measures. Let's consider how each of these changes bears on emotive doxastic embeddings.

Moss doesn't supply a semantics for embeddings under attitude verbs, but there is a natural way for her to treat embeddings under attitudes of acceptance, which we've already encountered in our discussion of Yalcin above: take attitudes of acceptance to have the informational structure of a probability measure, so that one accepts $\text{MOD}\phi$ just in case one's acceptance state is a member of the set of measures denoted by $\text{MOD}\phi$ (cf. Rothschild (2012)). However, a now familiar problem arises for extending this account to cover embeddings under emotive doxastics: these attitudes are at least partly preferential, and so their informational structure does not seem to be that of a probability measure. As a result, there is trouble generating any non-trivial truth-conditions for the attitude reports on a view like Moss's that takes the semantic value of an epistemically modalized sentence to determine a set of probability measures.

One way out of this problem is to integrate the view of Anand & Hacquard (2013) discussed above: take the modalized complement to characterize doxastic structure, and the non-modalized prejacent to characterize preferential structure, in a hybrid state. In order to block the entailment from hopes $\text{MOD}\phi$ to believes $\text{MOD}\phi$, Moss could maintain that the doxastic component of an emotive doxastic is a state of believing-possible, so that hoping $\text{MOD}\phi$

²⁴ Moss (2015) §2.2.

involves believing $\text{might-MOD}\phi$, where the nested modals are non-redundant. The resulting semantics retains some of the virtues of Anand & Hacquard's proposal (e.g. explaining why hoping $\text{MOD}\phi$ typically involves hoping ϕ), but loses others (e.g. the explanation of the resistance of *must* to embed no longer goes through). The biggest problem, however, is that the account would still validate the bad entailment from hopes $\text{MOD}\phi$ to hopes ϕ . We could avoid this problem if the modal had some effect in characterizing the preferential component of the emotive doxastic attitude. But as we noted above, it is puzzling how some bit of probabilistic content can characterize an attitude that seems to lack probabilistic structure. And if it could, the complexities of something like Anand & Hacquard's semantics would probably become superfluous. After all, the motivations for splitting the characterizing roles of modalized complements on the one hand, and modal prejacent on the other, were derived from an inability of modalized complements to interact with preferential components of attitudes. Once we allow that it is possible for modalized complements to interact with preferential components, we seem to lose motivations for the idea that there's any interesting 'split' interaction in the ascriptions to begin with.²⁵

In short, though the appeal to non-redundant modals makes some progress for the expressivist, we are still led back to the key issue of how a modal that characterizes probabilistic structure could also characterize a preferential state, which seems to lack that structure. Indeed, this is especially pressing on Moss's semantics since all modalized constructions express sets of probability measures. Since we are consistently driven to this problem, it is time to see how an expressivist might meet it head on, by dropping the hybrid semantics and taking a view on which epistemically modalized sentences supply values that can characterize both attitudes of acceptance and preferential attitudes.

Since modals express measures or sets of measures, the natural way to ensure they can characterize preferential states is to foist something like proba-

²⁵ If modals characterize preferential attitudes, Anand & Hacquard also lose their explanation of another puzzling datum: namely, epistemics' resistance to embedding under so-called desideratives—attitude verbs like 'want' or 'desire'. Anand & Hacquard argue that these embeddings are marked since desideratives have no doxastic component, and their preferential component cannot be characterized by an epistemic modal. We, by contrast, take the failure of the entailment from hopes $\text{MOD}\phi$ to hopes ϕ to demonstrate that modals do characterize preferential attitudes, in which case the infelicity of epistemics under desideratives remains an outstanding question. We'll have more to say on the matter in n. 44 below.

bilistic structure on those states—for example, by having the preferential states themselves be modeled as a measure, a set of measures, or as a relative ranking not merely on propositions but on measures. But though this allows the compositional machinery of modals to engage with that of preferential attitude verbs, it raises a pressing question about the interpretation of the formalism: what is a state of mind that is modeled as a set of (somehow ‘preferred’) measures, or as a (somehow ‘preferential’) ranking of measures, supposed to be? What is it like to be in a preferential state of mind in which some measure is in the state, or in which a measure is highly ranked in it?

This state of mind can’t be a matter of preferring *that* one’s credal state be a member of a given set of measures or preferring *that* one’s credal set be a subset of this set. Preferences like these are preferences toward propositions *about* one’s credences. But hoping that Alice might come to your party isn’t a matter of preferring that you’re in a particular type of doxastic state. The problem here is one that expressivists themselves press forcefully against many descriptivist semantics for modals: namely, that those semantics implausibly make attitudes towards modalized contents higher-order (see [Yalcin \(2007, 2011\)](#), [Rothschild \(2012\)](#), and [Moss \(2013\)](#)). Moreover, treating preferences towards sets of measures in this way threatens to collapse what was supposed to be an expressivist semantics into a local form of descriptivism. The semantics makes modalized hopes comprise preferences that a single, particular truth-conditional proposition hold. It just happens to be a proposition about one’s credal states satisfying certain properties. Expressivists might maintain that this concession to descriptivism only holds for modals embedded under emotive doxastics. But aside from appearing *ad hoc*, this reply leaves expressivists with the problem discussed in §2.1 of needing to give what is effectively a descriptivist-friendly account of the infelicity of embedding epistemic contradictions under emotive doxastics which, if successful, would undermine a key motivation for embracing expressivism in the first place. Note that neither of the foregoing problems can be helped by changing whose credences one’s hopes are about.

But if being in a preferential state that is modeled by a measure, a set of measures, or a ranking of measures, doesn’t amount to having preferences that credal states satisfy certain conditions, then how are they to be understood? A comparison to a non-preferential attitude like belief is instructive. We can make sense of what it is to have one’s doxastic state characterized by a set of measures—where this is not equivalent to having a full belief in some propo-

sition about that set of measures—since doxastic attitudes come in degrees that are fruitfully modeled via a probability measure. This suggests that the proper way to have one's attitude characterized by a set of measures is for the attitude itself to have graded structure. But as we've already stressed in §2.1, although preferences *do* have some kind of graded structure—the degree to which one prefers or disprefers—these degrees are clearly not what modal constructions characterize. It seems, then, that the only way for expressivists to make sense of the state of mind characterized by a set of measures is to embrace the radical thesis that a preferential attitude like hope (or the preferential component of this attitude, if we take hope to be a hybrid state of mind) has some other, *sui generis* graded structure.

This is a bitter pill to swallow. Expressivism about epistemic modality is plausible in part because the states of mind it adverts to—credences, sets of credal states, etc.—are fruitfully deployed in independent theorizing about the nature of doxastic states, especially in formal epistemology. But there is no such precedent for the proposed expressivist account of modalized hope: probabilistically graded preferential attitudes do not figure in independent theorizing about the nature of preferential attitudes (for example, in Decision Theory). Indeed, it's not even clear what a probabilistically graded preferential attitude would amount to. As we've already argued, that structure cannot be the degree to which one hopes or fears. And, as just noted, it cannot collapse into structure one prefers or disprefers to exist in mental states, or in the world, without becoming a local form of descriptivism that undermines some of the more important cases for expressivism. Expressivists seemingly must claim that the probabilistic structure of an emotive doxastic state is *sui generis*, and unexplained. But hopes and fears don't seem obscure in this way: their preferential structure appears unremarkable and adequately representable by something like a ranking on propositions.

Existing expressivist views are inevitably stuck with the foundational challenge of making sense of probabilistically graded preferential states, and motivating their relevance to the study of preferential attitudes if they are to have any hope of properly capturing embeddings of epistemic modals under emotive doxastics. Since we cannot see how this key problem can be addressed, we provisionally set expressivist views aside.

3 CONTEXTUALISM

Before proceeding, let's review some desiderata that we've uncovered in our discussion of expressivist views. Letting B stand for *believes* and E for an emotive doxastic like *hopes*, we have:

- (a) $B(\Diamond\phi)$ tends to accompany agnosticism whereas $E(\Diamond\phi)$ tends *not* to accompany preferential indifference.
- (b) $E(\phi \wedge \text{MOD}\neg\phi)$ sounds infelicitous.
- (c) $E(\text{MOD}\phi) \not\models B(\text{MOD}\phi)$ (and, plausibly, $E(\text{MOD}\phi) \models \neg B(\text{MOD}\phi)$).
- (d) $E(\text{MOD}\phi) \not\models E(\phi)$. Relatedly, MOD is not vacuous in $E(\text{MOD}\phi)$.
- (e) $E(\text{MOD}\phi)$ is compatible with, and typically held alongside $E(\phi)$.
- (f) \Diamond in $E(\Diamond\phi)$ seems to hedge the doxastic commitments of the attitude holder.
- (g) $E(\text{MOD}\phi)$ doesn't seem to report a 'higher-order' attitude about preferential or doxastic states.
- (h) $E(\Box\phi)$ tends to sound infelicitous.

We take (a)–(h) to substantiate the claim that emotive doxastics have an interesting truth-conditional and inferential behavior that promises to teach us important lessons about the semantics of modals. Indeed, we've learned that many of the most prominent expressivist treatments of modals have important and challenging work to do if they want to capture the data. We now turn to see whether the descriptivist fares any better.

3.1 MENTALISTIC AND VERITIC CONTEXTUALISM

We saw that a key component of the challenge that emotive doxastics raise for expressivists is a compositional 'mismatch' problem. Expressivists use modals scoped under attitudes to characterize structural properties of doxastic states that are expressed in bare modal assertions. But porting this semantics to emotive doxastics is in danger of creating a compositional crash if emotive doxastic states sometimes lack relevant structural features (e.g., probabilistic structure) or have a wholly different structure (e.g., that of a preference ranking rather

than an information state). An obvious virtue of a descriptivist position is that it need never face such worries. For the descriptivist, modalized constructions just express ordinary truth-conditions. Such truth-conditions can interact compositionally with emotive doxastics, no matter what the structure of an emotive doxastic state. So the most serious problem for the expressivist simply doesn't arise for the descriptivist.

But the descriptivist is hardly without challenges of her own. Those challenges are just relocated to the task of getting proper truth-conditions for the embeddings, including doxastic and preferential attitudes alike. Even the former has proven to be a vexing issue for descriptivists, as we alluded to in our discussion of Yalcin (2007). And some of the existing issues for the descriptivist are exacerbated when we consider emotive doxastic embeddings.

To see why, let's focus on what has some right to be called the orthodox descriptivist position. On this view, epistemic modals descriptively characterize bodies of information contained in the mental states, like states of belief or knowledge, of some contextually determined agents *S*. Asserting *it might be that p* is to claim that *p* is (for example) compatible with the information in the relevant states.²⁶ This is a simple statement of fact—a statement whose truth turns on how things stand with the information in the relevant mental states of the members of *S*.

If we take such a view and its simplest integration with attitude ascription, *A hopes might ϕ* will be equivalent to *A hopes that it is compatible with what the members of S believe/know that ϕ* . These are unacceptable truth-conditions, regardless of who is used to populate the relevant group *S*.

Suppose Sam hopes that there might soon be a female president of the United States. And suppose Sam exists at a bizarre world *w* such that everyone actual or in any close worlds is agnostic, for epistemically irrational reasons, about virtually everything pertinent to upcoming elections. This is intuitively not a description of the kind of outcome Sam is hoping for. Learning this would not, for example, be a way for Sam to learn that his hopes are fulfilled. But this is a situation where it is compatible with the beliefs and knowledge of *any* group of individuals *S* that there will soon be a female president.

Such examples can easily be multiplied, and extended to cover views which take even 'accessible knowledge' to settle epistemic possibility or likelihood.

²⁶ See, e.g., DeRose (1991) and Stanley (2005) for descriptivist views appealing to group and individual mental states respectively.

You can fear that species in remote depths of the oceans might be dying off due to climate change without having your fears realized simply by the fact that no one is in a position to know how climate change affects these creatures.²⁷ You could, as the last surviving member of the human race, hope that the achievements of mankind are likely to leave their mark on the universe somehow, without hoping anyone did, will, or would have a high credence in this event taking place.

The examples reveal a general pattern. When you are concerned with the possibility or likelihood of some ϕ , you do not necessarily have a concern for *how individuals do, or will, think* about ϕ . Rather, you seem to have something closer to a concern for ϕ itself. Note the problem is especially acute if we take the relevant set of individuals S for the embedded modal to be settled by the embedding attitude—that is, for S to contain only the attitude holder at the time the attitude is held.²⁸ Hoping something might be the case, or is likely, then becomes a truly bizarre hope about what one's own current hopes or credences are like. But one's hopes, in our cases of interest, are obviously not confirmed or disconfirmed by introspection.

These problems should sound familiar. They are essentially the problems with 'higher-orderism' about attitudes that Yalcin (2007, 2011) has raised for descriptivists, and we raised for certain extensions of expressivism in §2.3. Beliefs about what is possible or likely do not seem to be beliefs about various doxastic states. Yalcin concluded from this problem that embedded modals should function to characterize a property of an embedding attitude state, distinct from that of the state's relating to any particular proposition.

What we're seeing is that the problems of higher-orderism are worse for descriptivists when emotive doxastics do the embedding, and (with the results of §2.3) that the problems are actually not helpfully resolved by taking the modals to characterize the embedding attitude (whether directly, as in a domain se-

²⁷ Intuitively, what one fears here is that certain unknown truths leave open or make likely the extinction of these creatures. The example is constructed to make these truths not easily knowable: the truths might concern the ecology of inaccessible ocean locations. Thus, it's hard to see how a knowledge-based semantics for modals could include these truths in the modal base without having the modal base fixed by the contextually relevant *possible* knowledge. But talk of merely possible knowledge is roughly co-extensive with talk of what is true at a world. And so what appears to be a knowledge-based semantics turns out to be just a different way of stating the view we go on to defend below, according to which modal domains are fixed by the contextually relevant truths.

²⁸ Cf. Dorr & Hawthorne (2013) pp.906–7.

mantics, or indirectly, by making the attitude a higher-order state directed to propositional content about itself).²⁹

We think the data here pressures the descriptivist to consider an alternative to the orthodox contextualist position. All the problems above stem from the fact that modals characterized a contextually supplied set of mental states. What we need to consider are *non-mentalistic* forms of contextualism, according to which modals scoped under attitudes (and hence plausibly unembedded as well) characterize non-mentalistic bodies of information. The expressivist obviously should not be comfortable with this suggestion. After all, the expressivist is constitutively committed to unembedded modals functioning to characterize classes of mental states—that’s how bare modals can help ‘express’ them. But the descriptivist is free to explore alternatives.

The remaining task of this section is to outline a non-mentalistic form of contextualism and explore its implications for the attitudinal embeddings. The important parts of our positive proposal are compatible with many different compositional implementations of contextualism, but for specificity we’ll work within a framework similar to that developed in Kratzer (1977, 1981, 1991). On this view, epistemic modals are effectively quantifiers over possible worlds: they describe how things stand with the worlds left open by some relevant body of information (understood here as a set of propositions). One way to settle the relevant body of information is to use a so-called restrictor phrase. Consider: *according to the report, there may be life on Mars*. Here the phrase *according to the report* marks the information contained in the report as the relevant information characterized by the modal, and the sentence describes this information as compatible with life on Mars: there exists some world compatible with the truth of all the relevant propositions at which there is life on Mars. When a sentence contains no restrictor phrase, the relevant body of information is settled by the conversational context. In either case, Kratzer models the selection of the relevant information by a so-called modal base, f . f is a function that maps a world to a set of propositions—those constituting the relevant information at that world. So for example, *according to the report* expresses a function from a world w to the set of propositions that specify the information contained in

²⁹ It’s worth noting that the problems we’ve raised can’t be avoided merely by denying that mental state concepts figure directly in modalized contents (cf. Dorr & Hawthorne (2013) p.907 n.58). For instance, taking Sam’s hopes to be about the propositions that happen to be known or believed by the group still makes it too easy for him to get what he hoped for if the group is irrational.

the report at w . Modals then quantify over the worlds in the intersection of these propositions:

$$\llbracket \Diamond \phi \rrbracket^{c,w} \text{ is true iff } \exists w' \in \bigcap f^c(w) : \llbracket \phi \rrbracket^{c,w'} \text{ is true}$$

That is, *might* ϕ is true at a context-index pair just in case the contextually relevant body of information at the world w of the index is compatible with ϕ . We will adopt this compositional semantics for possibility and necessity modals (the latter differ only in expressing universal quantification). We will also use a modal base in our semantics for probability modals:

$$\llbracket \Delta \phi \rrbracket^{c,w} \text{ is true iff } Pr(\llbracket \phi \rrbracket^c \mid \bigcap f^c(w)) > \frac{1}{2}$$

On this semantics, *probably* ϕ is true at a context-index pair just in case the probability of ϕ is greater than .5 given the contextually relevant body of information at the world w of the index (we'll say more about the function Pr shortly).

And the following thesis encapsulates the minimal lessons that we draw from the emotive doxastic embeddings.

Veritic Contextualism: For epistemic modals MOD without explicit restrictors, MOD by default quantifies over a domain of possible worlds $\bigcap f^c(w)$ restricted by all of the contextually relevant true propositions at w .

To get a feel for how this view works, consider a sentence like *John might have cancer*. According to *Veritic Contextualism*, context selects a modal base f^c whose value at w consists in a set of contextually relevant true propositions at w . The modal base in this example might be a function from a world w to the set of true propositions describing John's symptoms at w and describing how these symptoms covary with cancer at w .

What distinguishes *Veritic Contextualism* from mentalistic forms of contextualism is that the latter maintains, while the former denies, that the propositions restricting the modal domain are universally, or by default determined as a function of what is known, believed, or otherwise related to the mental states of some contextually relevant subject or group. On *Veritic Contextualism*, epistemic modals tell us what is possible, necessary, or likely in view of the facts or truths.³⁰

³⁰ Kratzer has also recently stressed the dangers of pervasively linking modal bases to mental

Consider again the modal base from our example sentence above—namely the function from every world w to the set of true propositions describing John’s symptoms at w and how these symptoms covary with cancer at w . It may be true that the output of this function at the world of the context of use consists in a set of propositions known by the speaker of the context of use—that is, it may be true that the speaker knows John’s symptoms and how they covary with cancer at the actual world. But this does not mean that the modal base delivers a set of propositions $f^c(w)$ known by the speaker of c at all worlds w . Indeed, it is likely that at some worlds, the speaker fails to know what John’s symptoms are and how they covary with cancer. This is the sense in which *Veritic Contextualism* is a non-mentalistic form of contextualism: unlike orthodox contextualism, *Veritic Contextualism* does not hold that f^c tracks the information contained in mental states by making $f^c(w)$ return a set of propositions known by the contextually relevant individual or group for all worlds w . Rather, f^c tracks the relevant truths by returning the set of true propositions describing John’s symptoms at w and how these symptoms covary with cancer at w for all worlds w . The result is that the proposition expressed by the modalized sentence concerns the relevant truths, not some body of knowledge: the truth of the sentence at a world turns on how things stand with the non-mentalistic facts at that world, not on what is known by the speaker or some group at that world.

However, it is worth flagging that none of this implies that knowledge need have no role to play in a theory of epistemic modality. But any such role will typically be pragmatic, not semantic: the fact that certain propositions are known to some individual or group may well explain *why* a non-mentalistic property of propositions plays the role of epistemic modal base.

One example of this phenomenon arises from the following type of case: sometimes the only reason a modal happens to characterize a particular set of truths is that all of these truths are known to some individual or group at the context of utterance. For instance, if we are investigating a murder, and I say *Jones might be the murderer*, the relevant truths might be the ones concerning the kind of information that is known to me—e.g. the fingerprints and the

states (Kratzer (2012) p.50, p.98) and goes so far as to say that “epistemic modals’ do not have any necessary connection to knowledge” (*ibid.* p.50). One point Kratzer is emphasizing is that we shouldn’t expect epistemic modal bases to always be mentalistic. We, by contrast, want to go a step further and claim that (absent restrictors) epistemic modals are *typically* non-mentalistic.

eye witness testimony. And these truths might be the relevant ones precisely because they are known to me. But again, it does not follow that the modal thereby describes knowledge, in the sense of being true at a world depending on how things stand with what I know at that world. Rather, in this example, the proposition is true at a world depending on how things stand with the facts about the fingerprints and the eye witness testimony at that world.

Not only is it consistent with *Veritic Contextualism* that mentalistic facts may determine non-mentalistic modal bases, but it is even consistent with the claim that mentalistic properties sometimes play the role of modal base itself. The way this could happen is rather simple. *Veritic Contextualism* claims that modals quantify over worlds compatible with salient or relevant truths. What is salient or relevant can obviously depend on the interests of speakers. And one thing that may sometimes be of interest to speakers is what is compatible with, or probabilified by, what is known to one or more agents. In such cases, what counts as salient may precisely be what some set of agents knows. In such cases, and only in them, facts about knowledge may enter into an explanation of the nature of the modal base itself. Still, the veritic contextualist maintains that ‘being known to such-and-such agents’ is one of innumerable many features that could be salient or relevant to agents using modals and is, accordingly, liable to be atypical.

An analogy with ordinary quantifiers may help illustrate both of the foregoing points. It is obviously incorrect to claim that ordinary quantifiers like *all* or *some*, as a matter of their received semantics, quantify only (or even by default) over objects agents know, or are aware of, or believe to exist, and so on. The correct account of the semantics of quantifiers is that they quantify over salient or relevant objects satisfying a quantifier restrictor, with no reference to mental states.

But, of course, everyone should allow that the fact that we are aware of some objects can play a role in explaining why a given quantifier ranges over those objects. Sometimes it may be that a certain class of objects becomes salient to a speaker and hearers precisely because those agents are currently aware of objects in that class. This is analogous to the first point above: that sometimes mentalistic facts can make non-mentalistic modal bases salient, without anything mentalistic being hard-wired into the semantics for modals.

Also, all should allow that what is salient to speakers using ordinary quantifiers depends on the interests of those speakers, and sometimes such speakers

may be interested in a set of objects known to one or more agents. In such cases, and only those, the fact that an object is known to some agents may figure more directly in an explanation of why it is in a quantifier's domain: the domain of the quantifier may vary world by world with the objects satisfying the relevant mentalistic property. Again, the fact that this occurs does not mean that there is any pressure to treat quantifiers as having some reference to states of knowledge or belief hard-wired into their semantics. Indeed, even cursory investigation of uses of quantifiers reveals that a standard of being-known is rarely what determines a quantifier's domain in this direct way. So setting up the semantics for quantifiers to make this their default behavior would constitute a serious mistake. The veritic contextualist claims that mentalistic contextualists have made an analogous error with epistemic modals, and that the interaction of epistemic modals with emotive doxastics is especially helpful at revealing it.

So, in general terms, *Veritic Contextualism* maintains that *might* ϕ is true at a context-world pair $\langle c, w \rangle$ just in case the truths at w satisfying standards of relevance at c leave open ϕ . The truth of *likely* ϕ turns on features of a probability measure Pr , which encapsulates relations of evidential support.³¹ Thus, *likely* ϕ is true at a context-world pair just in case the contextually relevant truths at w support ϕ more than not- ϕ . Since modals state a worldly condition by targeting truths, there's no need to complicate our semantics for attitude reports.

$$\llbracket A \text{ believes } \phi \rrbracket^{c,w} \text{ is true iff } \forall w' \in \mathcal{B}_A^w : \llbracket \phi \rrbracket^{c,w'} \text{ is true.}$$

That is, we retain the original Hintikka semantics for belief, on which A believes that ϕ just in case ϕ is true at all of the worlds compatible with A 's beliefs. This leads to the following treatment of attitudes embedding epistemics.

$$\begin{aligned} \llbracket A \text{ believes } \Diamond \phi \rrbracket^{c,w} \text{ is true} & \quad \text{iff} \quad \forall w' \in \mathcal{B}_A^w : \llbracket \Diamond \phi \rrbracket^{c,w'} \text{ is true} \\ & \quad \text{iff} \quad \forall w' \in \mathcal{B}_A^w : \exists w'' \in \bigcap f^c(w') : \llbracket \phi \rrbracket^{c,w''} \text{ is true} \end{aligned}$$

To believe ϕ is possible is to believe the world is such that the relevant truths don't rule out ϕ . To suppose ϕ is likely is to suppose that the world is such that the relevant truths disproportionately favor ϕ . Similarly, to hope that ϕ

³¹ The idea that evidential relations can be modeled by a probability measure is familiar from Bayesian epistemology and discussions of so-called epistemic or evidential probability (see, for instance, Williamson (2000) Ch.10). But we don't want to take a stand here on how the value of Pr is fixed for a given use of a modal. We discuss two alternatives in §4.

is possible or likely is simply to hope things are such that the relevant truths leave open, or favor, ϕ . Modals characterize not our beliefs and hopes, considered as information states, but batches of truths that would obtain if what was believed or hoped were true. For example, one way for one's hopes that it's likely to rain tomorrow to be satisfied is for a large storm front to be moving towards the area at roughly the right rate. Note that if context and the relation of evidential support (settled by *Pr*) are fixed, hoping ϕ is likely, and believing ϕ is likely, involve distinct attitudes to the same proposition—distinct ways of relating to the same ways the world might be.

Obviously this proposal is abstract and skeletal. We haven't said anything yet about *which* truths get privileged in a particular context, for a particular token modal. This is clearly an important and extremely complex issue. But it is not our aim here to address it. One reason for setting the issue aside is that we are not sure that there are simple, informative explanations of modal domains of quantification, any more than there are for more ordinary domains of quantification which are sensitive to contextual cues in myriad ways.³² Moreover, even if there are such explanations, there is no reason to believe that veritic contextualism would be prevented from taking them on. To take one example, veritic contextualism is entirely consistent with the proposal of Dowell (2011) that contextual relevance is determined by the speaker's publicly manifestable intentions. The only difference with her semantics is that we would treat the speaker's intentions as being unconstrained by a mentalistic modal base.

Most importantly, though, our theory's virtues reviewed below are largely independent of one's final choice of contextual privileging. So let's consider how a veritic semantics interacts with semantics for emotive doxastics, and for attitudes more generally.

3.2 RECONSIDERING THE DATA

At the beginning of §3, we noted eight key pieces of data ((a)–(h)), and we've now seen that expressivists and mentalistic contextualists are not obviously able to account for all the data at once. The dialectical situation is complex: sometimes problems handled well by the expressivist present strong challenges to the mentalistic contextualist, and sometimes the reverse holds. Here, we want to focus on our positive story, arguing that veritic contextualism accounts for all the data at once in a simple and natural way. It is important to bear in

³² Cf. Glanzberg (2007).

mind that for each point we treat, some rival view like mentalistic contextualism or expressivism may handle the data. They may even handle it in roughly the same way. Our main claim is not that existing theories cannot account for any data point as well as our view. It is rather that no existing rival theory is able to capture all the data, as ours can.

- (a) $B(\Diamond\phi)$ tends to accompany agnosticism whereas $E(\Diamond\phi)$ tends *not* to accompany preferential indifference.
- (d) $E(\text{MOD}\phi) \not\models E(\phi)$. Relatedly, MOD is not vacuous in $E(\text{MOD}\phi)$.
- (e) $E(\text{MOD}\phi)$ is compatible with, and typically held alongside $E(\phi)$.

Embedded *might* helpfully conveys attitudinal indifference for belief and other attitudes of acceptance, but not with hope or fear. A veritic semantics has a simple explanation of why this would occur. Deliberation that leads to agnosticism about ϕ will generally entrain acceptance that relevant truths are compatible with ϕ and with its negation. In other words, if you've found that you haven't yet settled the issue of whether to accept ϕ , even after having considered the relevant truths, then you are liable to think the truths in question are compatible with ϕ (and its negation).

By contrast, deliberation that results in preferential indifference with regards to ϕ won't generally entrain preferences regarding the bearing of relevant truths on ϕ . In fact, it will generally entrain the opposite—preferential indifference to those truths. If you couldn't care less whether ϕ , you probably also couldn't care less whether the relevant facts leave ϕ open, or support ϕ .

Contrariwise, preferences for truths that leave open, or support ϕ tend to come alongside (more primary) preferences for ϕ . This is especially strong if we take on customary semantics for preferential attitudes involving contextually determined preference orderings. If you prefer the relevant truths leave ϕ open (over their not leaving it open) or that the relevant truths probabilify ϕ (over their not probabilifying it) then typically, but not always, this is precisely because you prefer ϕ (over not- ϕ). This explains a second, converse asymmetry between emotive doxastics and acceptance attitudes that embed modals: believing $\text{MOD}\phi$ doesn't tend to involve believing ϕ , but hoping or fearing $\text{MOD}\phi$ tends to involve hoping or fearing ϕ , as just noted. Moreover, the veritic contextualist can explain the strong connections here without having to posit anything as strong as an entailment, which we saw was problematic

for Anand & Hacquard's proposal. To take one of our previous examples: A student can hope a snowstorm tomorrow is likely by hoping that strong evidence for a snowstorm is about, without hoping there is a snowstorm. The student might do this on the presupposition that relevant evidence will be detected, and school proactively closed—where the proactive closure has more to do with the evidence, and less to do with whether there is a snowstorm or not.

Expressivist views had the hardest time with this data. Mentalistic contextualists may be able to take on board an analog to the explanation we give here, but in doing so they run headlong into problems with higher-orderism, which we consider again below.

(b) $E(\phi \wedge \text{MOD} \neg \phi)$ sounds infelicitous.

A veritic semantics takes epistemic modals to quantify over worlds compatible with 'the truths' at a world of evaluation (restricted by features of contextual relevance). As such *Veritic Contextualism* should be informed by the following pragmatic effect that draws on parallels between modals and quantifiers.

Salience. Epistemic modals function as quantifiers over worlds compatible with propositions true at a world, where those propositions meet contextual standards of relevance. So the value of a modal base at a world should be responsive to propositions true in that world, if salient as such, just as ordinary quantifier domains at a world are responsive to objects saliently satisfying a quantifier restrictor at that world. (In particular, for all modals MOD, if context c renders the proposition expressed by ϕ salient as true at w , typically $\llbracket \phi \rrbracket^c \in f^c(w)$.)

Salience embodies the recognition of a simple parallel between a veritic contextualist semantics for modals and treatments of ordinary quantifier domains. Quantified constructions like *all F s* and *some F s* select some contextually determined F s, to be described by material in quantifier matrices. The story of how context affects which objects are selected is, as already noted, incredibly complex. But it is generally agreed that something like the salience of certain objects, especially as provided by linguistic context (e.g. referring to a particular object) has a strong influence on whether that object figures in the con-

textually determined quantifier domain.³³ Modals are, on the proposal we've adopted, something like quantifiers which by default range over worlds constrained by contextually determined true propositions at a world. So embracing *Salience* is effectively to claim that modals behave like quantifiers vis-à-vis issues of salience.³⁴

The assumption of *Salience* is important, because it tells us how veritic semanticists should view epistemic contradictions like (24'): they should effectively be treated as instances of a more general kind of construction that we'll call a *quantified contradiction*. (24) is an example.

(24) # Mary likes coffee and no one likes coffee.

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

(24) is a conjunction with a quantifier in the second conjunct. The assertion makes mention of a person in its first conjunct. It thus puts forward the kind of entity over which the quantifier in the second conjunct ranges. Accordingly, the ordinary influence of salience ensures that Mary is picked up by the subsequent quantifier domain. The result is that assertions of (24) state contradictions in the context in which they're uttered: they claim that Mary likes coffee, but that no one (including Mary) likes coffee. This doesn't mean that (24) is contradictory in every context, but only that it will be interpreted as such without very strong countervailing contextual cues.

If a veritic semantics is governed by *Salience*, then sentences like (24') should tend to state contradictions for analogous reasons. (24') is a conjunction with a quantifier in the second conjunct: *might*. The assertion puts forward a proposition in its first conjunct as a truth at the actual world. It thus puts forward the kind of thing which enters the modal base of that quantifier at actuality: the truth of there being rain. Accordingly, (24') states that it's raining, and that

³³ See, e.g., von Stechow (1994), Stanley & Szabo (2000).

³⁴ It's important to emphasize that we do not put forward *Salience* as a complete account of how context affects modal domains. There may be cases in which a true proposition restricts a modal domain even though the proposition is not salient as true (consider the example Dorr & Hawthorne (2013) (p.879) use to motivate what they call 'constrained' readings of modals). The same goes for quantifier domains more generally: the domain of *all* may contain objects that do not saliently satisfy a restrictor (cf. Ichikawa (2011) p.384). However, none of these complications bear on the examples we discuss below. The key point for our purposes is the following sufficiency claim: if a proposition is salient as true at a world, it will typically restrict a modal domain at that world.

it's compatible with the salient truths (including the fact that it's raining) that it's not raining.³⁵

If we want to further understand the behavior of epistemic contradictions, we need only consider the corresponding behavior of quantified contradictions. To that end, let's consider some important features of the latter.

First, note that although (24) tends to sound like a contradiction, we are not committed to a general entailment from (25) to the negation of (26).

(25) No one likes coffee.

(26) Mary likes coffee.

No one likes coffee in context *c* only entails that the *c*-relevant individuals don't like coffee. These may or may not include Mary. Indeed for most ordinary contexts *c*, she will not be among the *c*-relevant individuals.

Second, note that the influence of salience has some retroactive force.

(27) ? No one likes coffee and Mary likes coffee.

(27) continues to sound problematic. But it is certainly an improvement over (24). And it can be improved further with contrastives like *but* or *however*.

Third, consider 'split' embeddings of quantified contradictions.

(28) Mary likes coffee and Pia thinks no one likes coffee.

³⁵ In coming to appreciate this aspect of our view, it may be helpful to contrast epistemic and circumstantial modals. On a veritic semantics, epistemics have a certain kinship with circumstantial modals which are customarily treated as also having modal bases determined by truths or facts. For example, *Petunias can grow here* may mean that Petunias could grow given the facts about the pH balance of the soil. The key difference between epistemics and circumstantials on our view is in the role of context in settling a modal base. Epistemics act (very roughly) as quantifiers explicitly restricted only by 'truths'—with context delimiting which truths are relevant. A circumstantial by contrast behaves like a quantifier explicitly restricted by 'truths of type *T*'—for example those facts relevant to ascertaining something like a disposition or capacity. In other words, we might think of a circumstantial as functioning (very roughly) as does a construction like "Given the truths of type *T*, ..." with epistemics functioning more like "Given the truths, ..." (where of course not all truths can be taken into account, and context will help delimit which are relevant). If a truth isn't of the right type *T*, even raising it to salience will have no effect on the truth-conditions of the first kind of construction (which is why *Petunias aren't growing here, though they could grow here* can be felicitous), but should have a very strong, though defeasible, effect on the truth-conditions of the second kind of construction.

It's a little tricky to say whether, in (28), Mary is in the quantifier domain of *no one* (though there is a definite sense Pia believes incorrectly). But there's no infelicity in the pronouncement. Suppose, for example, someone utters (28) at a party. Pia may know Mary likes coffee but (say) be unaware she's at the party. If so, (28) could intuitively be true, even though Mary likes coffee at all worlds compatible with what Pia believes. The rough account of this behavior should be relatively clear: the embedded quantifier *no one* is functioning to tell us about the properties of persons who satisfy conditions of salience and relevance *at Pia's belief worlds* (and perhaps even relevant to new standards of salience and relevance set in part by Pia's doxastic state, or our purposes in discussing it). Mary may or may not be among such persons, whether or not we mention her explicitly in a discussion which makes it clear that she *actually* satisfies conditions of salience and relevance.

But problems arise when we scope the first conjunct back under the attitude.

(29) # Pia thinks Mary likes coffee and no one likes coffee.

(30) # Suppose Mary likes coffee and no one likes coffee.

In these cases, the first embedded conjunct ensures Mary satisfies conditions of salience and relevance at the worlds described as compatible with Pia's beliefs, or our supposition worlds. We accordingly get exactly as strong a sense that a contradictory belief is attributed to Pia in (29) as we do for thinking that (24) states a contradiction. We also hear the instruction in (30) as incoherent.

Fourth, while it is a commonplace that quantifier domains include salient, relevant objects in their domains, this view is in no way committed to the claim that we can paraphrase quantificational talk using ordinary language words like *relevant* or *salient*. Compare (31) with (24).

(31) Mary likes coffee and no relevant/salient person likes coffee.

(32) Pia thinks that Mary likes coffee and no relevant/salient person likes coffee.

Whereas (24) sounds like a straightforward contradiction, (31) tends to sound better. Indeed, it is much easier to think of contexts in which (31) sounds perfectly appropriate, and true, even though (24) continues to sound aberrant. (32) can also sound basically fine.

There are several candidate explanations of this phenomenon. One is that the notions of salience and relevance in the linguistic theory of quantification are technical ones. The terms *salient* or *relevant* may be ambiguous or polysemous. Even if not, the technical uses of linguists may effectively introduce such ambiguity or polysemy. If so, when faced with the explicit uses in (31), charity may lead listeners to substitute less technical interpretations of the polysemous or ambiguous words, which no longer tend to generate contradiction.

We take no stand on whether this, or some other explanation, holds. We merely note that if quantifier domains are sensitive to salience or relevance, *some* such explanation must be available to explain the contrast between (24)/(29) and (31)/(32). So the standard treatment of quantifier domains as responsive to salience and relevance, in ways that tend to make quantified contradictions genuinely contradictory, isn't in any way committed to the view that we can adequately paraphrase quantified constructions using natural language terms. On the contrary, it is committed to coming up with some explanation of why this *isn't* the case.

These four points help us to understand the behavior of epistemic contradictions, which exhibit perfectly parallel behavior. Though, on our account, assertions of epistemic contradictions tend to state literal contradictions, this in no way commits the view to a general entailment from (25') to the negation of (26').

(25') It might not be raining.

(26') It's raining.

Indeed, most ordinary contexts *c* in which (25') is asserted will obviously be ones where the fact that it is raining (or, more generally, the answer to the question of whether it's raining) is not considered to be among the *c*-relevant truths.

As with (27), we find that salience works retroactively within a sentence, so that (27') continues to sound odd, though less so than (24') (as noted by Sorensen (2009), Willer (2013), among others).³⁶

³⁶ We noted that (27) could be improved not only by reversing the order of a quantified contradiction but using a contrastive, adding suitable linguistic context, or raising to salience other objects to figure in a quantifier domain. Jointly, these can effectively eradicate the infelicity. The same seems to be true of epistemic contradictions, as has also been noted by Sorensen (2009), Dorr & Hawthorne (2013), Moss (2015). See also Yanovich (2014) for a related discussion of felicitous epistemic contradictions.

(27') ? It might not be raining and it's raining.

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

(28') does not necessarily seem to report Pia's beliefs about what is compatible or incompatible with the fact that it's raining (though Pia seems to believe incorrectly).

(28') It's raining and Pia thinks it might not be raining.

So, as before, we seem to find that a modal scoped under the belief ascription is being used to tell us what relevant truths *Pia* thinks obtain—not the relevant truths we are claiming to obtain. Again it is only when we scope the first conjunct of the epistemic contradiction back under the attitude that we ensure that the proposition that it's raining figures in the modal base of *might*. For example, in (29') the proposition that it is raining is now put forward in the first conjunct as a truth at Pia's belief worlds, and raised to salience as such. And the embedded modal is telling us about compatibility with salient and relevant truths at those worlds. Thus we get exactly as strong a sense that Pia believes incoherently in (29') as we get that (24') states an incoherence. Relatedly, we tend to hear (30') as an incoherent instruction.

(29') # Pia thinks it's raining and it might not be raining.

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

Finally, if the parallel between epistemic modality and quantification holds good, we should expect (31')/(32') to behave like (31)/(32)—that is, they should sound significantly improved. And they do.

(31') It's raining and the relevant/salient truths leave open that it might not be raining.

(32') Pia thinks it's raining and the relevant/salient truths leaves open that it might not be raining.

Again, it doesn't matter for our account why it is that (31) is an improvement over (24). As long as that explanation doesn't appeal to a highly unusual expressivist treatment of quantification generally, we will have a perfectly acceptable parallel descriptivist account of why (31') improves over (24'). The cases

merely reinforce the parallel behavior between quantified and epistemic contradictions.

This final point introduces a very important caveat for our view: our semantics for *it might be that* ϕ is not safely tested with intuitions about the ordinary English paraphrases like *the relevant truths leave open* ϕ . The embedding behavior of the two foregoing sentences must be different given the contrast between (31)/(32) and (24)/(29). Talk of ‘relevant truths’ in giving and testing our proposed semantics is fine, as long as each of these caveats are borne in mind.

To get clearer on the semantics of embedded epistemic contradictions, we have the following.

$$\begin{aligned} \llbracket A \text{ believes } (\phi \wedge \Diamond \neg \phi) \rrbracket^{c,w} \text{ is true} & \quad \text{iff} \quad \forall w' \in \mathcal{B}_A^w : \llbracket \phi \wedge \Diamond \neg \phi \rrbracket^{c,w'} \text{ is true} \\ & \quad \text{iff} \quad \forall w' \in \mathcal{B}_A^w : \llbracket \phi \rrbracket^{c,w'} \text{ is true and} \\ & \quad \exists w'' \in \bigcap f^c(w') : \llbracket \neg \phi \rrbracket^{c,w''} \text{ is true} \end{aligned}$$

The belief report raises the proposition expressed by ϕ to salience as true at all of the worlds the report describes as compatible with A ’s beliefs in c . So *Salience* ensures that, for all such worlds w' , $\llbracket \phi \rrbracket^c \in f^c(w')$. But this means that there are no worlds in $\bigcap f^c(w')$ at which $\llbracket \phi \rrbracket^c$ is false. Thus, there will be no worlds w' such as to satisfy the second condition above. In effect, the attitude report describes A as believing an ordinary contradiction: A believes that ϕ and that the relevant truths (including ϕ) are compatible with $\neg \phi$.

Note that this same story can be told for hope and fear, whether those states have an informational or preferential structure, and even if we embed a probabilistic modal instead: the embedded propositions will always impose unsatisfiable conditions.³⁷ In fact, our theory has the immediate prediction that epistemic contradictions will tend to sound contradictory wherever they are embedded, including in conditional antecedents (Yalcin (2007)), as dis-

³⁷ We noted in n.12 that it is an open question whether preferential attitudes are subject to the same consistency norms as, say, belief, though it at the very least seems marked to attribute hope or fear in a bit of contradictory content (as in *A hopes* $\phi \wedge \neg \phi$), as opposed to a set of contents that are jointly inconsistent. *I hope Emeline arrives soon and I also hope she doesn’t* may seem to some like a coherent expression of ambivalence. But *I hope Emeline both does and doesn’t arrive soon* sounds confusing in just the way that (12)–(13) do, as our account would predict. Moreover, as we’re grateful to [blinded for review] for pointing out, this account is arguably an improvement over that of Anand & Hacquard (2013), which more controversially requires it to be problematic to hope ϕ and hope $\neg \phi$ (see n.14).

juncts (Dorr & Hawthorne (2013)) or scoped under other epistemics (Moss (2015)). Most importantly, for present purposes, our theory can pair freely with any theory of attitude reports.

The reason for all this, as alluded to above in our discussion of attitude reports, is that embedded conjunctions continue to make the proposition expressed by their first conjuncts salient as true at all the worlds the conjunction rules in, whether those worlds are being ruled in as believed, or supposed, or are being ruled in as one among many ways the world might be in a disjunction. By *Salience*, the first proposition in a conjunction will restrict modal domains at all worlds where the whole conjunction is true. This occurs for epistemic contradictions roughly as it does for quantified contradictions. Even when embedded in a disjunction, a conditional antecedent, or an attitude report, the conjunction given by a quantified contradiction makes the object mentioned in the first conjunct salient as satisfying the restrictor at all worlds ruled in by the conjunction as a whole. This is why the quantified contradiction embodies a contradictory claim wherever it is embedded.

This explanation of epistemic contradictions is very simple. But it is worth stressing that the simplicity doesn't come for free. It is only afforded if one embraces *Veritic Contextualism*. If modals habitually characterize the structure of mental states, *Salience* is much harder to motivate, and the question immediately arises as to why there is almost no temptation to get a coherent reading of embedded epistemic contradictions on which they characterize a speaker's mental state.

To appreciate this point, it may be helpful to contrast our explanation with a recent attempt by Dorr & Hawthorne (2013) to capture the same data within the general framework of the orthodox semantics for modals. Dorr & Hawthorne claim that in addition to the normal uses of *might* on which it expresses compatibility with the information in knowledge states, *might* has a special 'constrained' interpretation, on which the modal base is expanded with answers to a question or issue of whether ϕ (p.883). So *might* ϕ , constrained by the issue of whether ψ , is true at a world w just in case:

There is a world w' that is accurate with regard to whether ψ in w , and that is compatible with what is known in w , and ϕ is true at w' .

Linguistic context can make salient a way in which a possible world can be accurate, so that subsequent modals 'inherit' this constraint. For example, the

first conjunct in an epistemic contradiction like ϕ and *might not* ϕ renders salient a world's potential accuracy with respect to whether ϕ . So, provided the modal in the second conjunct receives a special constrained interpretation, this conjunct is true at the world of evaluation just in case there is a not- ϕ world that is both compatible with what is known and accurate with respect to whether ϕ in the world of evaluation. However, the first conjunct is true at the world of evaluation just in case there is no such world, and thus the sentence as a whole expresses a contradiction.

This so far explains why, *if* a modal receives the hypothesized constrained reading in an epistemic contradiction, we will get an infelicity. But since epistemic contradictions are almost always marked, including in various embeddings, we need to know why modals in epistemic contradictions tend to get constrained readings. Dorr & Hawthorne supplement their theory with an elaborate pragmatics to explain why speakers are drawn to these contradictory, constrained interpretations, as opposed to more charitable interpretations on which the modal simply characterizes the contextually relevant knowledge.³⁸

The success of Dorr & Hawthorne's account seems to hinge on the correctness of their complex, and controversial, pragmatic account. Our explanation, by contrast, relies only on *Salience*—a principle motivated by the general pragmatics of quantifier domain restriction combined with the assumption that modals have a veritic semantics. The pragmatics invoked is completely familiar, and the explanation of all epistemic contradictions is uniform and simple. This significant gain in simplicity seems to us to correspond to a commensurate gain in plausibility. The only way for Dorr & Hawthorne to match the simplicity of our explanation would be to claim that hereditarily constrained readings are default interpretations of modals—an option they consider but don't pursue.³⁹ This seems to us to be done with good reason: it doesn't seem plausible that such readings could be defaults. Recall that constrained read-

³⁸ To give a flavor of the complexity: Embedding failures in conditional antecedents and disjunctions are explained by appeal to a default assumption of epistemic transparency (that is, lack of ignorance of what one knows) along with implicatures allegedly generated by connectives. Together these drive us to constrained readings. This raises the questions: Why do the implicatures go through, rather than being cancelled precisely because they would otherwise result in contradictory interpretations? And even if transparency can't be avoided, why don't we instead charitably search for unconstrained readings to avoid the contradiction? This is patched with a further pragmatic thesis: there is a default preference for explicit rather than tacit reference to contextual parameters in these cases (pp.899ff.).

³⁹ Dorr & Hawthorne (2013) p. 888.

ings of modals are sensitive to answers to a question or issue somehow made salient, often linguistically. If constrained readings were default interpretations, it would seem that merely raising the issue or question of whether ψ should virtually force a constrained reading of $\Diamond\phi$, rendering salient a world's potential accuracy with respect to whether ψ . But sentences like (33) and (34) seem to show this doesn't occur.

(33) Is it raining? Well, it might be raining.

(34) I wonder whether it's raining. Well, it might be raining.

If constrained readings of modals, which after all are constrained by a salient issue or question, were default readings, then the question raised explicitly in the first sentence, and implicitly in the second, seems like it should raise to salience a possible world's accuracy with respect to the question of whether it is raining. But if the modal in the second sentence inherits this constraint, *it might be raining* becomes equivalent to *it is raining*. Surely this is not the default reading of the modal—indeed, it's not clear this is even a possible reading. As far as we can see, the best way to avoid this problem (at least, provided one wants to maintain constrained readings are default readings) is to drop sensitivity to an issue or question, and instead build in sensitivity to something more like a salient truth, as advocated on our view. Then raising a question wouldn't suffice to alter the modal base, even if the new form of 'constrained' readings were made defaults. Of course, if one both takes modal bases to be influenced by salient truths, and takes these interpretations to be defaults, one ends up with our view.⁴⁰

To sum up: no existing expressivist or mentalistic contextualist view is able to account for the epistemic contradiction data in its full generality, especially given the data from emotive doxastic embeddings. But on the veritic view, all

⁴⁰ It's worth noting that the drive to the complex pragmatics of Dorr & Hawthorne arises because they posit two types of interpretation. The hard question is why charity doesn't lead us to pick the interpretation that creates a felicitous or informative reading. As they note (pp.887–8), this is actually *just* as serious a problem for someone like Yalcin—who allows that modals can receive a descriptivist reinterpretation. On our view, by contrast, there is really a single *kind* of interpretation of epistemic modals (that can of course vary in interpretation with context), and it is governed by the same pragmatic constraints that govern all uses of quantifiers. This gives us a ready-made answer to the question of why we don't seek out any rival contextual contributions to avoid contradictions. Given the parallel between epistemic and quantified contradictions, whatever explains why charity doesn't lead us to felicitous readings of the latter will also carry over to the former.

epistemic contradictions are neatly explained as instances of the broader phenomenon of quantified contradiction.

(c) $E(\text{MOD}\phi) \not\models B(\text{MOD}\phi)$ (and, plausibly, $E(\text{MOD}\phi) \models \neg B(\text{MOD}\phi)$).

(f) \Diamond in $E(\Diamond\phi)$ seems to hedge the doxastic commitments of the attitude holder.

In a way, the explanation of the blocked entailment for us is simple: hoping a state of affairs obtains generally involves not believing it obtains. $\text{MOD}\phi$ expresses a (contextually-dependent) state of affairs on our view. So hoping $\text{MOD}\phi$ should not involve believing $\text{MOD}\phi$.

But for reasons noted earlier, the story needs to be a little more complex. It is plausible not only that hoping ϕ comes with lack of belief in ϕ , but also belief that ϕ is epistemically possible. If so, as we noted in §2.3, avoiding the entailment from *hopes* $\text{MOD}\phi$ to *believes* $\text{MOD}\phi$ may also require nested modals to be non-redundant. Fortunately, a veritic semantics has a natural explanation of why modals are sometimes non-redundant. Context selects veritic bases modal by modal. Since modals characterize sets of truths, nested modals characterize truths about truths. As such, context can naturally select different modal bases for pairs of nested modals, by effecting some interesting separation of truths about truths about ϕ from truths about ϕ .

We may not generally make fine distinctions between such truths, but this seems to be roughly what we do in cases where we witness blocked entailments from hoping $\Diamond\phi$ to believing $\Diamond\phi$. Consider our earlier example.

- (15) (a) I hope we might get back together.
 (b) I believe we might get back together.

If I assert (15a), and hope entails believing-possible, then I should think it's an open possibility that my ex and I might get back together. This seems right. But I'm also not prepared to assert that we might get back together—that seems too strong. According to the present account of default modal targets, the belief that it's an open possibility that my ex and I might get back together is a belief that the truths bearing on truths bearing on whether we'll get back together leave open that the truths bearing on whether we'll get back together are compatible with us getting back together. There's a loose, but intuitive sense

of how and why we might separate these different batches of truths. I might count as truths pertinent to our reunion things like my ex saying she's been thinking about it positively or negatively. I might count as truths about truths things like that her friends haven't reported to me that she's thinking about our reunion positively or negatively. If I am separating these kinds of truths, they can be associated with different modals in the obvious way.

Note that if nested possibility modals aren't redundant, this seems to occur because the nested modals weaken one's commitments. Saying *it's possible that Brad Pitt might be coming to your party* seems to involve a weaker commitment to the possibility of Pitt's attendance than saying *Brad Pitt might be coming to your party*, to the extent these say anything different at all.⁴¹ This is captured on our view because absent any additional constraints on modals, we have the following (using subscripts to track modal bases associated with tokens).⁴²

Weakening. When context selects modal bases making $\Diamond_1\Diamond_2$ non-redundant (i.e., not equivalent to \Diamond_2), then $\Diamond_1\Diamond_2\phi$ is logically weaker than $\Diamond_2\phi$. That is:

$$\Diamond_2\phi \models \Diamond_1\Diamond_2\phi \text{ and } \Diamond_1\Diamond_2\phi \not\models \Diamond_2\phi$$

The first entailment follows trivially from the fact that *Veritic Contextualism* takes only true propositions to restrict modal domains. The failure of the second entailment is then required by non-redundancy.

The resulting account allows us to give a satisfying explanation of the semantic relationship between hoping ϕ and hoping $\Diamond\phi$. First, the account blocks the entailment from *hopes* *MOD* ϕ to *believes* *MOD* ϕ for the reasons just given. Second, the account frees us up to claim that hoping *MOD* ϕ entails, presupposes, or generally requires *not* believing *MOD* ϕ . If hopes generally require agnosticism, then because the modalized complements now have truth-conditions, there is no reason why agnosticism cannot extend to what it expresses, especially if the truth-conditions of the complement are stronger than those of nested modal constructions.

⁴¹ Wolf (2017) also notes this weakening effect.

⁴² Allowing intrasentential shifts in modal bases may obviously require some amendments to the semantics. One option is to attach indices to modal tokens in syntactic structure; another is to lean on whatever mechanism will be used to handle intrasentential contextual shifts in general—for example to account for contextual shifts in repeated uses of demonstratives (see Moss (2015) and Braun (1996), respectively).

Third, the account explains the sense in which the function of a modal in *hopes* $\Diamond\phi$ is predominantly to hedge the doxastic commitments of the hoper. We've already noted that, except in rare cases, hopes that ϕ is possible come with more primary hopes for ϕ itself.⁴³ This means that *A hopes* $\Diamond\phi$ doesn't report a significant difference from *A hopes* ϕ in *A*'s preferential state. But even so, the former does signal a significant difference in *A*'s doxastic state. If hope requires agnosticism about what is hoped for, and hoping $\Diamond\phi$ does not entail believing $\Diamond\phi$, then hoping $\Diamond\phi$ involves believing $\Diamond\Diamond\phi$, where the nested modals are non-redundant. As we've noted, any account that captures plausible features of non-redundant possibility modals will generate modal weakening, and therewith weaken the doxastic commitments of the relevant hoper. Our account predicts that this weakening, whatever it amounts to, constitutes the intuitive hedging that the modal in *hopes* $\Diamond\phi$ creates. Indeed, given what we've said about the preferential similarities between *A hopes* $\Diamond\phi$ and *A hopes* ϕ , our account predicts that the typical function of embedded possibility modals in emotive doxastic reports is to weaken the doxastic commitments of the attitude holder.⁴⁴

The foregoing virtues of the veritic contextualist view largely owe to its descriptivist character. As a result, it may be possible for the mentalistic contex-

⁴³ It might be objected that our account now predicts that hoping $\Diamond\phi$ is incompatible with hoping ϕ : hoping ϕ requires believing $\Diamond\phi$, while hoping $\Diamond\phi$ requires agnosticism about $\Diamond\phi$. But this ignores that the 'implicit' modal associated with *hopes* ϕ can involve a different modal base than the explicit modal in *hopes* $\Diamond\phi$, in which case these hopes can again be compatible. Indeed, blocking the entailment from hoping $\Diamond\phi$ to believing $\Diamond\phi$ requires precisely this kind of separation between explicit and implicit modals.

⁴⁴ The foregoing account also gives us insight into two further issues: the fact (acknowledged in n.8) that some embeddings of epistemics under emotive doxastics may require contextual set-up, and the fact that epistemics tend to resist embeddings under desideratives. Embeddings under emotive doxastics may sometimes need contextual set-up if that set-up is needed to make sufficiently clear how distinct contextual contributions are made to modal bases to ensure non-redundancy. The explanation of the resistance of desideratives to embed epistemics is also broadly pragmatic. Since the attitudes reported by desideratives have no doxastic component, embedding possibility modals under these attitude verbs generates no such doxastic hedge. And since one typically prefers $\text{MOD}\phi$ only if one prefers ϕ , epistemics under desideratives add little to the characterization of one's preferential attitudes. As a result, when preferences are directed to the modal preadjacent, there is no real function for the embedded modal to play in a desiderative report, leading such a report to be degraded. Of course, there may be special contexts in which one only or primarily has preferences for how the truths bear on ϕ rather than for ϕ itself (recall the examples of blocked entailments from *hopes* $\text{MOD}\phi$ to *hopes* ϕ). Such contexts should license epistemics under desideratives on the current proposal. But, intriguingly, such contexts do seem to license the desiderative embeddings, as Anand & Hacquard (2013, pp. 22–23) themselves note.

tualist to replicate these results, though their explanation would have to differ in some details that we won't explore here.

- (g) $E(\text{MOD}\phi)$ doesn't seem to report a 'higher-order' attitude about preferential, credal, or other doxastic states.

In our recent discussion of 'orthodox' descriptivism, we noted that such views may implausibly make modalized hopes and fears higher-order attitudes. We also saw this was a danger for certain elaborations of the role of probabilistic content in preferential attitude ascriptions in §2.3. We noted that Yalcin (2007, 2011) has pressed related problems for orthodox descriptivists, even for modals embedded under attitudes of acceptance. To claim that *A* believes it might rain is in danger of claiming that *A* believes that it is compatible with what *A* believes that it rains. Such a view seems to rule out attributions like (35) to animals that lack the capacity for higher-order cognition.

- (35) My dog Fido thinks you might give him a bone.

The problems for standard contextualist views arise because the content they attribute to modal claims is mentalistic. For example, a typical mentalist contextualist semantics for *might* looks something roughly like the following:

$\llbracket \Diamond \phi \rrbracket^{c,w}$ is true iff ϕ is compatible with what the speaker of *c* knows at *w*.

The truth conditions of the content expressed by *might* ϕ at a context of utterance vary world by world with what some agent knows. Accordingly, to say that someone believes or hopes ϕ might happen invariably attributes to their attitude a sensitivity to how things stand with some mental states. This leads to all the aforementioned problems. And, as we noted before, expressivists are in danger of generating similar problems depending on how they try to cope with probabilistic embeddings.

A veritic semantics is able to skirt these difficulties. A veritic semantics looks roughly as follows:

$\llbracket \Diamond \phi \rrbracket^{c,w}$ is true iff ϕ is compatible with the truths at *w* satisfying standards of salience/relevance at *c*.

Such a semantics does not raise any of the problems we've seen arise for expressivists or orthodox descriptivists. Considering batches of true propositions merely involves considering ordinary worldly information. This could involve information from testimony, surveys, or test results, but also much simpler information gleaned from ordinary perception.

The information relevant to the interpretation of modals is, by stipulation, not constrained by the information in any mental states. And because this worldly information can come from any other sources, being in touch with it requires no special cognitive sophistication. Fido, for example, can believe or hope you might give him a bone, or that you are likely to give him a bone, because Fido can be sensitive to indications that you're about to give him one (knowing looks, familiar smells, your possession of a bone, your being in a good mood, e.g.), and believe or hope that those indications are about.

One might worry that the veritic contextualist has problems of its own, in that it attributes awareness of standards of relevance or salience to believers and hopers—facts that may be even more sophisticated than simple facts about knowledge.⁴⁵ But this worry is unfounded. Standards of salience and relevance in the veritic view play the role they always do: in helping to determine the content expressed by an expression type at a context. They are not themselves part of the content so-expressed. This is reflected in the fact that on the veritic semantics above, the truth-value of *might* ϕ does not vary, world (of evaluation) by world (of evaluation) with facts about salience and relevance at those worlds. Another way of putting this, using the terms introduced by Kaplan (1989), is that facts about salience and relevance belong to the character, and not the content, of a modalized construction on the veritic view. Accordingly, no awareness of, or responsiveness to, such standards are required by thinkers said to believe or hope modalized complements. The problem with standard mentalistic views is that they make mentalistic information part of the content of such complements.

Mentalistic contextualists are the views most immediately affected by problems of higher-orderism. But the way that expressivists try to avoid those problems—by having modals directly characterize embedding attitudes—is what leads to the most serious problems from emotive doxastic embeddings. As such, veritic contextualists are in a uniquely advantageous position to account for the relevant data in its full generality.

⁴⁵ We're grateful to an anonymous reviewer for raising this objection.

- (h) $E(\Box\phi)$ tends to sound infelicitous.

Treating the foregoing data from emotive doxastic embeddings required little more than adopting a veritic semantics. By contrast, accounting for the infelicity of embeddings of *must* under emotive doxastics may require more controversial commitments. Here we will limit ourselves to two observations that illuminate the phenomenon and help explain why we think plausible accounts of the data will be compatible with a veritic semantics.

The first observation is that *must*—unlike possibility and probability modals—resists embedding under possibility modals.⁴⁶ Compare:

- (36) It's *possible* that the report *might* mislead.
 (37) The report *could* be *likely* to mislead.
 (38) ? It's *possible* that the report *must* be misleading.

The second observation is that on a veritic semantics, we can salvage the view that hoping/fearing requires believing epistemically possible. If so, hoping/fearing *must- ϕ* would require believing it is epistemically possible that *must- ϕ* .

These two observations together strongly suggest a kinship between the infelicity of *might-must- ϕ* , and embeddings of *must* under emotive doxastics. The latter entail or presuppose the felicitous combination of the very sequence of nested modals which, it seems, is problematic. Thus, possibility and probability modals felicitously embed under emotive doxastics, while *must* does not:

- (39) We fear the report *might* mislead.
 (40) We fear the report is *likely* to mislead.
 (41) ? We fear the report *must* be misleading.

So it seems like the problems with nested modals are primary, and so the problems from emotive doxastic embeddings should be explicable in terms of the infelicity of the nested modal construction.

To be clear: we doubt that necessity modals invariably resist embedding under possibility modals, but we also doubt that *must* invariably resists embedding under emotive doxastics. Instead, our suggestion is that *must* very

⁴⁶ Thráinsson & Vikner (1995) make the same observation with respect to epistemic modals in Scandinavian languages.

often resists embedding under possibility modals and that the (in)felicity of $\Diamond\Box\phi$ at a context c patterns with the (in)felicity of hopes/fears $\Box\phi$ at c .

The source of the infelicity in $\Diamond\Box\phi$ is one we won't speculate on, since we suspect that it involves more controversial claims about the syntax, semantics, or pragmatics of nested modal constructions than we'd like to take on here.⁴⁷ What is important to recognize is that the challenge of explaining the resistance to embedding is a problem for all semantics for modals. And a veritic semantics is a good candidate to explore our options for coping with all the data, since it has the virtue of enabling some nested modals to function non-redundantly, and it seems to capture the crucial observation that ties hopes/fears- $\Box\phi$ to $\Diamond\Box\phi$. (It is of course not necessarily unique in this respect: any other view that captures these two bits of data is equally well poised to explain the data.) Since we only mean to explore a skeletal veritic semantics here, we don't want to make more commitments on the behavior of nested modals for now. So we leave this as an issue for further research.

4 BROADER LESSONS OF THE EMBEDDINGS

We've argued that a skeletal veritic framework captures almost the entire data set ((a)–(g)) with minimal constraints on which truths are considered salient or relevant in a context. Though accounting for the remaining bit of data ((h)) seems to require further commitments, the data seems helpfully illuminated by the semantics. If the veritic framework is on the right track, what lessons can we take away for expressivist and descriptivist treatments of epistemics more generally?

As we saw in §2, existing expressivist views all face serious challenges in accommodating embeddings under emotive doxastics. But we believe there may be a way of combining veritic and expressivist semantics into a novel, hybrid semantics which captures the foregoing data from embeddings under emotive doxastics while supporting a partially expressivist pragmatics of assertion, as well. If so, developing this sort of semantics may be the most promising line for expressivists to take.

The possibility of such a hybrid view stems from the fact that thinking about how things stand with a batch of truths really involves something like a pair of attitudes: beliefs about which truths obtain, and beliefs or presupposi-

⁴⁷ See [Butler \(2003\)](#) for a potential syntactic explanation.

tions about which propositions the batch of truths counts in favor of believing and to what extent. Disagreement seems possible not only over what is true, but over what the truths count in favor of believing. We chose to encapsulate information about evidential relations in a probability measure *Pr*. So if we can take different attitudes about what constitutes evidence for what, *Pr* should be variable, and perhaps shift its value when scoped under attitude constructions. How should this variability be encoded? The answer turns on our conception of evidential support.

One way to conceive of evidential support is in broadly factualist terms. There are many different ways to do this. Perhaps there are facts about what counts as evidence for what, and how much evidential support is given, which can be encoded in a probability measure. Perhaps there are more fundamental facts about objective probabilities that settle facts about evidential support. Perhaps the best way to understand evidential support is in some completely different, non-probabilistic way. If we take any of these options, the important point is that evidential relations become part of the fabric of the world. Accordingly, when we encode variability of evidential support, we should add a world parameter to our probability measure so that it can pick up its value from the ways a thinker believes, supposes, or hopes the world to be. The result will be a purely descriptivist semantics for epistemic modals.

But we may also try to make sense of evidential support in non-factualist terms, perhaps in connection with a position like subjective Bayesianism. Then our probability measure may encode information about credal states that aren't settled on a factual basis. Accordingly, it would be a mistake to encode the variability of evidential support with variability in the way the world might be. We should instead make amendments similar to those proposed by Yalcin. For example, we could add a shiftable information state parameter *s* into the index, allowing *Pr* to pick up its value from that parameter, and use attitude verbs to shift the parameter to the credal state of the attitude holder (say).

If we try to develop an expressivist semantics along the foregoing lines, what happens? It will turn out that all uses of modals will involve hybrids of worldly and non-worldly information—for example, in a domain semantics their semantic values would exhibit non-trivial variability along both the world and information state parameters. We could integrate such values into a pragmatics of assertion with a distinctly expressivist character. One example of such a hybrid expressivist view is the following: a use of $\Delta\phi$ expresses the speaker's

high conditional credence in ϕ given the relevant truths and constitutes a recommendation for conversational participants to coordinate on information states having this property. That is, listeners are to get into an information state in which their conditional credences and beliefs about the relevant truths line up in the following way: the relevant truths at the worlds compatible with one's information state are such that one's conditional credence in ϕ given these truths is high. To believe that ϕ is likely is just to be in an information state having this property. To hope that ϕ is likely is to hope that the world contains facts of the sort that would lead one to have a high credence in ϕ .

The descriptivist component of this theory is that the attitude holder's concern is for how things stand with the contextually relevant truths—the attitude holder essentially hopes for there to be truths that make ϕ likely. The expressivist component concerns what it *takes* for a set of truths to make a proposition likely: likelihood is settled by the attitude holder's conditional credences, not worldly facts about, say, objective probabilities.

This version of expressivism may well be capable of co-opting our descriptivist explanations of the embedding data: after all, on the proposed hybrid expressivist view, the objects of emotive doxastic attitudes are just *that the facts be a certain way*. What the other virtues and vices of such an expressivist position are is a matter for future research: to our knowledge, no forms of semantically hybrid expressivism have yet been explored in any detail.⁴⁸

Nonetheless, there are some very important senses in which expressivism about epistemic modals, and arguments for it, are constrained on a hybrid expressivist view. First, as just noted, the seemingly best candidates for expressivist discourse—discourse involving bare, unembedded modals—will at best still turn out to involve an admixture of worldly, factualist information and non-worldly information about (say) credal states. There is no *pure* expressivist discourse here. Second, it is this worldly information which has done all the work so far in resolving puzzles about the embedding behavior of epistemic modals. Every single virtuous feature of the theory we discussed in section three owed to its unambiguously descriptivist components. If the expressivist adopts a semantic hybrid theory, she will be giving up any motivations for expressivism from attitude reports.

What about descriptivists? As we have stressed, the veritic framework we've

⁴⁸ However, see [Schroeder \(2013\)](#) for discussion of an analogous view in metaethics (what we are calling “hybrid expressivism” corresponds to what he calls “relational expressivism”).

appealed to is skeletal. But we have also argued it is the best framework within which a descriptivist view can be developed. If so, obviously the next task is to explore the question of how context selects veritic modal bases for modals, especially with ‘bare’, unembedded uses of epistemics.

The task of saying how context selects veritic modal bases is obviously complex. It will involve revisiting issues, besides embeddings, that have been used to motivate the shift away from contextualist descriptivism, like retraction and eavesdropping phenomena. But although this is undoubtedly a complex task, we think the focus on a veritic view holds some promise of strengthening the descriptivist position. After all, when issues like retraction and eavesdropping are used to apply pressure to the contextualist descriptivist, they are typically used as part of a ‘process of elimination’ argument meant to rule out the most plausible descriptivist options. But typically the views eliminated are framed directly or indirectly in mentalistic terms—that is, in terms of what one or more agents do know, or could deduce from what they know, or what evidence is ‘easily available’ for them to know (e.g., Egan et al. (2005), MacFarlane (2014)). To adopt a veritic semantics is to reject such descriptivist views. We should no more appeal to mental states (even possible, or ‘nearby’ epistemic states) when describing which propositions form the modal base of epistemics than we should appeal to such states in stating what persons end up being quantified over by *everyone* or *someone*. As such, the shift to a veritic semantics raises concerns that some fundamental assumptions underlying the debates over the semantics of epistemics, often shared both by descriptivists and their opponents, should be abandoned.

REFERENCES

- P. Anand & V. Hacquard (2013). ‘Epistemics and Attitudes’. *Semantics and Pragmatics* 6(8):1–59. [1], [2], [3], [4], [10], [11], [13], [19], [39], [45]
- D. Bolinger (1968). ‘Post-posed Main Phrases: An English Rule for the Romance Subjunctive’. *Canadian Journal of Linguistics* 14:3–30. [10]
- L. Bovens (1999). ‘The Value of Hope’. *Philosophy and Phenomenological Research* 59(3):667–81. [12]
- D. Braun (1996). ‘Demonstratives and Their Linguistic Meanings’. *Noûs* 30(2):145–173. [44]
- J. Butler (2003). ‘A Minimalist Treatment of Modality’. *Lingua* 113(10):967–996. [49]

- K. DeRose (1991). 'Epistemic Possibility'. *Philosophical Review* 100:581–605. [24]
- K. DeRose (1998). 'Simple 'Might's, Indicative Possibilities and the Open Future'. *Philosophical Quarterly* 48(190):67–82. [2]
- C. Dorr & J. Hawthorne (2013). 'Embedding Epistemic Modals'. *Mind* 122(488):867–913. [25], [26], [34], [37], [40], [41]
- J. J. L. Dowell (2011). 'A Flexible Contextualist Account of Epistemic Modals'. *Philosophers' Imprint* 11(14):1–25. [31]
- A. Egan, et al. (2005). 'Epistemic Modals in Context'. In G. Preyer & G. Peter (eds.), *Contextualism in Philosophy*, pp. 131–170. Oxford University Press. [52]
- A. Egan & B. Weatherson (eds.) (2011). *Epistemic Modality*. Oxford University Press, Oxford. [55]
- A. Falaus (2010). 'Alternatives as sources of semantic dependency'. *Semantics and Linguistic Theory (SALT)* 20:257–306. [12]
- B. Geurts (2005). 'Entertaining alternatives: Disjunctions as modals'. *Natural Language Semantics* 13(4):383–410. [11]
- M. Glanzberg (2007). 'Context, Content, and Relativism'. *Philosophical Studies* 136(1):1–29. [31]
- I. Hacking (1967). 'Possibility'. *Philosophical Review* 76(2):143–168. [2]
- V. Hacquard (2006). *Aspects of Modality*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA. [1], [2]
- V. Hacquard (2010). 'On the event relativity of modal auxiliaries'. *Natural Language Semantics* 18(1):79. [1], [2]
- V. Hacquard (2011). 'Modality'. In K. von Heusinger & P. Portner (eds.), *Semantics: An International Handbook of Natural Language Meaning*. Mouton de Gruyter. [2]
- V. Hacquard & A. Wellwood (2012). 'Embedding Epistemic Modals in English: a Corpus-Based Study'. *Semantics and Pragmatics* 5:1–29. [1], [2], [3], [13]
- I. Heim (1992). 'Presupposition Projection and the Semantics of Attitude Verbs'. *Journal of Semantics* 9(3):183–221. [10]
- J. Hintikka (1969). 'Semantics for propositional attitudes'. In D. J. H. J. W. Davis & W. K. Wilson (eds.), *Philosophical logic*, pp. 21–45. Reidel, Dordrecht. [5]
- J. Ichikawa (2011). 'Quantifiers and Epistemic Contextualism'. *Philosophical Studies* 155(3):383–398. [34]

- D. Kaplan (1989). ‘Demonstratives’. In J. Almog, J. Perry, & H. Wettstein (eds.), *Themes From Kaplan*, pp. 481–563. Oxford University Press. [47]
- A. Kratzer (1977). ‘What *Must* and *Can* Must and Can Mean’. *Linguistics and Philosophy* 1(3):337–55. [26]
- A. Kratzer (1981). ‘The Notional Category of Modality’. In H.-J. Eikmeyer & Rieser (eds.), *Words, Worlds, and Contexts, new approaches to word semantics*, pp. 38–74. Walter de Gruyter, Berlin. [26]
- A. Kratzer (1991). ‘Modality’. In A. von Stechow & D. Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research*, pp. 539–50. Walter de Gruyter, Berlin. [26]
- A. Kratzer (2012). *Modals and Conditionals: New and Revised Perspectives*. Oxford University Press. [28]
- D. Lassiter (2011). *Measurement and Modality: The Scalar Basis of Modal Semantics*. Ph.D. thesis, New York University. [3]
- J. MacFarlane (2014). *Assessment Sensitivity*. Oxford University Press. [52]
- S. Moss (2013). ‘Epistemology Formalized’. *Philosophical Review* 122(1):1–43. [2], [14], [21]
- S. Moss (2015). ‘On the Semantics and Pragmatics of Epistemic Vocabulary’. *Semantics and Pragmatics* 8(5):1–81. [2], [4], [15], [18], [19], [37], [40], [44]
- P. Pettit (2004). ‘Hope and Its Place in the Mind’. *The ANNALS of the American Academy of Political and Social Science* 592:152–65. [12]
- P. Portner (1997). ‘The semantics of mood, complementation, and conversational force’. *Natural Language Semantics* 5(2):167–212. [12]
- D. Rothschild (2012). ‘Expressing Credences’. *Proceedings of the Aristotelian Society* 112(1pt.1):99–114. [2], [18], [19], [21]
- T. Scheffler (2008). *Semantic Operators in Different Dimensions*. Ph.D. thesis, University of Pennsylvania. [12]
- M. Schroeder (2013). ‘Tempered Expressivism’. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 8*. Oxford University Press. [51]
- R. Sorensen (2009). ‘Meta-agnosticism: Higher Order Epistemic Possibility’. *Mind* 118:777–84. [37]
- R. Stalnaker (1984). *Inquiry*. MIT Press, Cambridge, MA. [10]
- J. Stanley (2005). ‘Fallibilism and Concessive Knowledge Attributions’. *Analysis* 65(2):126–31. [24]
- J. Stanley & Z. G. Szabó (2000). ‘On Quantifier Domain Restriction’. *Mind and Language* 15(2&3):219–61. [34]

- E. Swanson (2011). ‘How Not to Theorize about the Language of Subjective Uncertainty’. In *Egan & Weatherson* (2011), pp. 249–269. [2], [3], [15], [16]
- H. Thráinsson & S. Vikner (1995). ‘Modals and Double Modals in the Scandinavian Languages’. *Working Papers in Scandinavian Syntax* 55:51–88. [48]
- H. Truckenbrodt (1997). ‘On the semantic motivation of syntactic verb movement to C in German’. *Theoretical Linguistics* 32:257–306. [12]
- F. Veltman (1996). ‘Defaults in Update Semantics’. *Journal of Philosophical Logic* 25(3):221–261. [6], [18]
- E. Villalta (2000). ‘Spanish Subjunctive Clauses Require Ordered Alternatives’. In *Proceedings of SALT X*. [10]
- E. Villalta (2008). ‘Mood and Gradability: An Investigation of the Subjunctive mood in Spanish’. *Linguistics and Philosophy* 31:476–522. [10]
- K. von Fintel (1994). ‘Restrictions on Quantifier Domains’. [34]
- M. Willer (2013). ‘Dynamics of Epistemic Modality’. *Philosophical Review* 122(1):45–92. [6], [18], [37]
- T. Williamson (2000). *Knowledge and its Limits*. Oxford University Press. [30]
- L. Wolf (2017). ‘Modal Concord is not Modal Concord’. *Annual Conference of the German Linguistic Society*. [44]
- S. Yalcin (2007). ‘Epistemic Modals’. *Mind* 116(464):983–1026. [1], [2], [4], [8], [21], [24], [25], [39], [46]
- S. Yalcin (2010). ‘Probability Operators’. *Philosophy Compass* 5(11):916–937. [2]
- S. Yalcin (2011). ‘Nonfactualism about Epistemic Modality’. In *Egan & Weatherson* (2011), pp. 295–332. [1], [2], [4], [6], [21], [25], [46]
- S. Yalcin (2012). ‘A Counterexample to Modus Tollens’. *Journal of Philosophical Logic* 41(6):1001–1024. [2]
- I. Yanovich (2014). ‘Standard Contextualism Strikes Back’. *Journal of Semantics* 31(1):67–114. [37]
- H. Zeijlstra (2007). ‘Modal Concord’. In T. Friedman & M. Gibson (eds.), *Proceedings of SALT 17*, pp. 317–332. Cornell University Press. [2]