

Embedding Epistemic Modals

CIAN DORR

University of Oxford

cian.dorr@philosophy.ox.ac.uk

JOHN HAWTHORNE

University of Oxford

john.hawthorne@philosophy.ox.ac.uk

Abstract:

Seth Yalcin has pointed out some puzzling facts about the behaviour of epistemic modals in certain embedded contexts. For example, conditionals that begin ‘If it is raining and it might not be raining, ...’ sound unacceptable, unlike conditionals that begin ‘If it is raining and I don’t know it, ...’. These facts pose a *prima facie* problem for an orthodox treatment of epistemic modals as expressing propositions about the knowledge of some contextually specified individual or group. This paper develops an explanation of the puzzling facts about embedding within an orthodox framework.

1. Yalcin’s data

There is an old-fashioned, ‘propositionalist’ account of epistemic modals on which ‘It might be the case that ϕ ’ and ‘It is possible that ϕ ’ are context-dependent sentences which express propositions about knowledge.¹ A particular use of ‘It might be the case that ϕ ’, whether bare or embedded, expresses a proposition roughly equivalent to those expressed by ‘ α does not know that not ϕ ’, or ‘ α is not in a position to know that not ϕ ’, where α refers to some relevant

¹ See, for example, Hacking 1967, Kratzer 2012 (Ch. 1), DeRose 1991, and Hawthorne 2007.

individual or group.² Seth Yalcin (2007) has introduced some new data that seem to challenge this equivalence. The data involve certain embedded occurrences of sentences like (1a) and (1b):

- (1) a. It is raining and it might not be raining
b. It is raining and I don't know that it is raining

There is no particularly notable contrast between (1a) and (1b) when uttered on their own: both are very odd. The oddity of the Moore-paradoxical (1b) can be explained as follows: in asserting the first conjunct, one is in some sense representing oneself as knowing that it is raining, and thereby committing oneself to something inconsistent with what one asserts with the second conjunct. The propositionalist can exploit this explanation to account for the oddity of (1a) by positing that in the absence of special contextual pointers, 'It might not be raining' will be interpreted as expressing that the *speaker* does not know (or is not in a position to know) that it is raining.³

But as Yalcin notes, the oddity of (1a) persists in various embedded constructions where (1b) sounds much better. Yalcin focuses on two particular contrasts of this sort:

² Read quotations including Greek variables as quasi-quotations.

³ In fact, all that is needed for this explanation is the claim that 'might' is by default indexed to a group including the speaker. If we think of groups as knowing whatever any of their members know, a speaker who knows that P can never belong to a group which fails to know that P. De Rose (1991) claims that the relevant group must *always* contain the speaker, so that 'Might ϕ and I know that not ϕ ' has *no* consistent reading. As others (e.g. Egan, Hawthorne, and Weatherson 2005) have pointed out, this seems too strong. When specific people or groups excluding the speaker are salient, 'might' can sometimes be interpreted as targeting their knowledge. (For examples see Sect. 3.2 below.) This raises the good question why, hearing (1b), we do not assume that the 'might' is being used to talk about some third party's knowledge, charitably passing over the Moore-paradoxical interpretations. Sect. 7 will introduce some considerations which suggest an answer.

- (2) a. # Suppose that it is raining and it might not be raining
- b. Suppose that it is raining and I don't know that it is raining
- (3) a. # If it is raining and it might not be raining, then ...
- b. If it is raining and I don't know that it is raining, then ...

Phenomenologically, at least when we encounter (2a) and (3a) out of the blue, we feel that we are being asked to entertain something like a contradiction. (2b) and (3b) by contrast sound fine. To account for this, Yalcin introduces some novel semantic machinery, broadly 'relativistic' in character, which we will briefly explain in section 2. Before accepting such an account, it is worth investigating whether the data can be explained within a familiar propositionalist framework, which attempts to explain facts about felicity by appeal to a semantic theory which specifies which propositions count as admissible interpretations of a given sentence, together with pragmatic factors which explain why some of the possible interpretations are easier to access than others in a particular conversational setting. We shall argue that, without straying too far from the old-fashioned account of epistemic modals, satisfactory explanations can be given for Yalcin's data as well as of related families of puzzles. We will mostly focus on the contrast between (3a) and (3b); in the final section we will return to (2a–b) along with other cases involving propositional attitude verbs.

2. Yalcin's explanation of the data

The key idea of Yalcin's theory is to replace the familiar form of semantics in which the semantic values of closed sentences are functions from possible worlds to truth values with one that adds an extra parameter. For Yalcin, the semantic value of a closed sentence is a

function which takes an ordered pair $\langle s, w \rangle$, where w is a possible world and s is a set of possible worlds (or ‘information state’), and yields a truth value.

The extra parameter is only relevant for sentences containing special expressions like ‘might’ and ‘if’. When a sentence contains none of these expressions, the question whether it is true at $\langle s, w \rangle$ depends only on w — such sentences are ‘information-insensitive’. ‘Not’ and ‘and’ are treated truth-functionally, and thus do not introduce any information-sensitivity:

‘Not ϕ ’ is true at $\langle s, w \rangle$ iff ϕ is not true at $\langle s, w \rangle$

‘ ϕ and ψ ’ is true at $\langle s, w \rangle$ iff ϕ and ψ are both true at $\langle s, w \rangle$

The information parameter is invoked in the obvious way in the semantics for ‘might’, which Yalcin treats as a sentential operator:

‘Might ϕ ’ is true at $\langle s, w \rangle$ iff ϕ is true at $\langle s, w' \rangle$ for some $w' \in s$

To explain Yalcin’s account of ‘if’, we will need a preliminary definition: a set of worlds s *accepts* a sentence ϕ iff ϕ is true at $\langle s, w \rangle$ for every $w \in s$. The proposal is that

‘If ϕ , ψ ’ is true at $\langle s, w \rangle$ iff the largest subset of s that accepts ϕ accepts ψ ⁴

⁴ What if s has no ‘largest’ subset that accepts ϕ ? For example, if s contains both worlds where some ordinary sentence ψ is true and worlds where it is not, ‘Not (might ψ and might not ψ)’ is accepted both by the subset of s containing only the worlds where ψ is false, and by the subset containing only the worlds where ψ is true, but not by s itself. Kolodny and MacFarlane (2010) suggest counting ‘If ϕ , ψ ’ as true at $\langle s, w \rangle$ iff *every* maximal subset of s that accepts ϕ accepts ψ . Another approach would count ‘If ϕ , ψ ’ as true at $\langle s, w \rangle$ iff ψ is

Note that on this account, whenever ‘Might ϕ ’ or ‘If ϕ , ψ ’ is true at $\langle s, w \rangle$, it is true at $\langle s, w' \rangle$ for every w' : such sentences are ‘world-insensitive’.⁵

If one defined a valid argument as one that preserves truth relative to each $\langle s, w \rangle$, we would have to say that modus ponens is invalid: if ϕ is an ordinary sentence false at w but true at some member of s , then relative to $\langle s, w \rangle$, ‘Not ϕ ’ and ‘If not ϕ , not might ϕ ’ are true while ‘Not might ϕ ’ is false. We get the same result if we only require valid arguments to preserve truth relative to all $\langle s, w \rangle$ meeting certain conditions — for example, that $w \in s$ — so long as there is some such $\langle s, w \rangle$ where the worlds in s do not all agree on the truth-values of ordinary sentences.⁶ However, Yalcin prefers a different account of validity as *informational consequence*, according to which an argument is valid when every set of worlds that accepts all the premisses accepts the conclusion. Modus ponens is valid by this criterion, but as Yalcin (2012b) points out, modus tollens is not, since any s containing both ϕ - and not- ϕ worlds will accept ‘Might ϕ ’ and ‘If not ϕ , not might ϕ ’ while failing to accept ‘Not ϕ ’. Also, various standard metarules like *reductio*, contraposition, and proof by cases fail for informational consequence.

Let us see how Yalcin’s semantics applies to (3a). First, note that the empty set accepts every sentence, and is the only set that accepts ‘ ϕ and it might be that not ϕ ’. (For s to accept

accepted by the *intersection* of all maximal ϕ -accepting subsets of s . The basic form of Yalcin’s explanation will work equally well on either of these approaches.

⁵ This will lead to some curious behaviour if we introduce ordinary non-epistemic modal operators in the obvious way. For example, if ‘It is metaphysically possible that ϕ ’ is true at $\langle s, w \rangle$ iff ϕ is true at $\langle s, w' \rangle$ for some metaphysically possible w' , then ‘It is metaphysically possible that ϕ ’ can be true at $\langle s, w \rangle$ even when ‘It is metaphysically possible that it might be that ϕ ’ is not.

⁶ Kolodny and MacFarlane (2010) adopt this kind of account of validity, and accept the conclusion that modus ponens is invalid. Their primary argument that some kind of logical revisionism is required turns on certain facts about the behaviour of deontic modals in conditionals. We think that these facts can be explained within a logically orthodox propositionalist framework, but we cannot go into details here.

‘It might be that not ϕ ’, it must either be empty or contain some w such that ‘Not ϕ ’ is true at $\langle s, w \rangle$; but in the latter case, s fails to accept ϕ .) Thus for any s , the null set is the largest subset of s that accepts ‘ ϕ and it might be that not ϕ ’. Since the null set also accepts ψ , ‘If ϕ and it might be that not ϕ , ψ ’ is true at every $\langle s, w \rangle$. The semantic values assigned to these sentences are thus the same as those assigned to conditionals with contradictory antecedents, like (4):

(4) If it is raining and it isn’t raining, then ...

This style of semantics can thus explain why (3a) generates the same kind of phenomenology as (4).

When people propose semantic theories that relativize truth for certain sentences to a parameter, they typically go on to say how values of that parameter get to be ‘operative’ on particular occasions in determining whether an utterance is felicitous. The same is true here — a full story will flesh out how particular information states enter into the felicity conditions for utterances of information-sensitive sentences. But we need not concern ourselves with this aspect of Yalcin’s theory. For whatever set of worlds we might consider in evaluating a given use of ‘If ϕ and might not ϕ , ψ ’, it will turn out to be defective in the way described.

What about the other contrast, with ‘suppose’? For this, Yalcin proposes a semantics based on the idea that for each person we can distinguish a particular set of worlds, those compatible with what the person is supposing (at a given time): ‘ α supposes that ϕ ’ is true relative to $\langle s, w \rangle$ iff the set of worlds compatible with *what the referent of α is supposing at w*

is one that accepts ϕ . Yalcin suggests that a parallel account will apply to many other attitude verbs, including ‘believe’, ‘suspect’, ‘think’, and ‘know’.⁷

Since only the empty set can accept ‘ ϕ and it might not be that ϕ ’, ‘ α supposes that (ϕ and it might not be the case that ϕ)’ can only be true at $\langle s, w \rangle$ if *nothing* is compatible with what the referent of α is supposing at w : it is just like ‘ α supposes that (ϕ and not ϕ)’. (2a) is thus an imperative to enter into this defective state of supposition. No wonder we balk at it.

Note that on this account, ‘ α supposes that ϕ ’ is information-insensitive: its truth value relative to $\langle s, w \rangle$ does not depend on the value of s . In this way the treatment of the attitudes follows the model of relativism or expressivism, as opposed to contextualism.⁸ Relativists who hold that ‘Marmite is tasty’ is true relative to some standards of taste and false relative to others do not say that ‘John MacFarlane believes that Marmite is tasty’ is relative in the same way.⁹ Similarly, expressivists who say that ‘Lying is wrong’ expresses a state of mind that is not the state of believing any proposition do not say that ‘Simon Blackburn believes that lying is wrong’ expresses such a state of mind. By contrast, attitude reports that embed normal context-sensitive sentences are themselves context-sensitive: ‘John MacFarlane believes that I like Marmite’ inherits the context-sensitivity of ‘I like Marmite’.

⁷ The application to ‘know’ yields especially problematic results. Since knowledge entails belief, any world compatible with one’s beliefs is compatible with one’s knowledge. Thus on Yalcin’s account, ‘ α knows that might ϕ ’ will follow from ‘ α believes that might ϕ ’. Worse still, since the actual world is compatible with everyone’s knowledge, ‘ α knows that might ϕ ’ follows from ϕ ! One could escape these problems by adding ad hoc additional clauses to the semantics for ‘know’, claiming that ‘ α knows that ϕ ’ is true at $\langle s, w \rangle$ iff (i) the set of worlds compatible with what α knows at w accepts ϕ , (ii) the set of worlds compatible with what α believes at w accepts ϕ , and (iii) ϕ is true at $\langle s, w \rangle$. But on this account, ‘ α knows that might ϕ ’ still follows from ‘Might ϕ and α believes that might ϕ ’, which seems objectionable: the usual reasons for thinking that true belief is not sufficient for knowledge do not suddenly lose their force when we are dealing with knowledge of how things might be.

⁸ Yalcin favours an expressivist interpretation of the machinery: see Yalcin 2011. We are not going to dwell on what, if anything, is at stake between expressivist and relativist glosses on the machinery.

⁹ Stephenson (2009) develops a version of relativism about both epistemic modals and predicates like ‘tasty’. For further characterization of relativism and extensive citations, see Cappelen and Hawthorne 2009.

3. Variants of Yalcin's semantics

3.1 Dynamic 'and'

A different way to invoke information-relative semantic values to explain Yalcin's data would replace Yalcin's truth-functional treatment of 'and' with the following 'dynamic' semantic clause:

'ϕ and ψ' is true at $\langle s, w \rangle$ iff ϕ is true at $\langle s, w \rangle$ and ψ is true at $\langle s[\phi], w \rangle$

where $s[\phi]$ is $\{w: w \in s \text{ and } \phi \text{ is true at } \langle s, w \rangle\}$. The resulting treatment of 'and' and 'might' is more or less equivalent to a standard dynamic semantics (von Stechow and Gillies 2007, Yalcin 2012a); the present formulation is due to Klinedinst and Rothschild (2012).¹⁰

The new clause for 'and' assigns different semantic values to 'ϕ and might not ϕ' and 'Might not ϕ and ϕ': the former is not true relative to any $\langle s, w \rangle$, whereas the latter is true relative to $\langle s, w \rangle$ provided that ϕ is true relative to $\langle s, w \rangle$ but false relative to $\langle s, w' \rangle$ for some $w' \in s$. Nevertheless, neither 'ϕ and might not ϕ' nor 'Might not ϕ and ϕ' is *accepted* by any nonempty set of worlds. So if we keep Yalcin's semantics for 'if', 'If ϕ and might not ϕ, ψ' and 'If might not ϕ and ϕ, ψ' will still be equivalent, since both will be true relative to all

¹⁰ In the standard presentation of the dynamic semantics, semantic values of sentences are functions from sets of worlds to sets of worlds. Each function v from ordered pairs $\langle s, w \rangle$ to truth values determines such a function f_v , given by $f_v(s) = \{w: w \in s \text{ and } v(\langle s, w \rangle) = \text{Truth}\}$. The determination is not quite one-one, since f_v and $f_{v'}$ will be identical provided that $v(\langle s, w \rangle) = v'(\langle s, w \rangle)$ for each s and $w \in s$. But it is not clear whether this matters, since no pair of sentences in the fragment treated by Yalcin will diverge only on those $\langle s, w \rangle$ for which $w \notin s$.

$\langle s, w \rangle$. However, the dynamic semantics for ‘and’ is naturally combined with the following alternative account of ‘if’ (cf. Gillies 2004):

‘If ϕ , ψ ’ is true at $\langle s, w \rangle$ iff $s[\phi]$ accepts ψ

When combined with the dynamic treatment of ‘and’, this predicts a difference in semantic value between ‘If ϕ and might not ϕ , ψ ’ and ‘If might not ϕ and ϕ , ψ ’. Since ‘ ϕ and might not ϕ ’ is not true relative to any $\langle s, w \rangle$, $s[\phi \text{ and might not } \phi]$ is always the empty set, so that ‘If ϕ and might not ϕ , ψ ’ is true relative to every $\langle s, w \rangle$.¹¹ However, $s[\text{‘Might not } \phi \text{ and } \phi\text{’}]$ need not be empty. So long as s contains some w such that ϕ is false relative to $\langle s, w \rangle$, $s[\text{‘Might not } \phi \text{ and } \phi\text{’}] = s[\phi]$, so if $s[\phi]$ fails to accept ψ , ‘If might not ϕ and ϕ , ψ ’ will be false relative to $\langle s, w \rangle$.¹²

Which of the two accounts fits the data better? The answer is not obvious. Changing the order of the conjuncts in the odd-sounding sentences does sometimes make for an improvement.¹³ For example, (5b) sounds better than (5a):

¹¹ Note that if we combined the revised clause for ‘if’ with the original, truth-functional account of ‘and’, even ‘If ϕ and might not ϕ , ψ ’ would no longer be vacuously true.

¹² Friends of dynamic semantics often favour an account of validity more demanding than Yalcin’s ‘informational consequence’, according to which ‘ ϕ_1, \dots, ϕ_n , therefore ψ ’ is valid iff for any s , $s[\phi_1] \dots [\phi_n]$ accepts ψ . This account fits in certain natural ways with the dynamic clauses for ‘and’ and ‘if’: for example, it entails that $\Gamma \vdash \text{‘If } \phi, \psi\text{’}$ whenever $\Gamma, \phi \vdash \psi$, and that $\phi, \psi \vdash \chi$ exactly when ‘ ϕ and ψ ’ $\vdash \chi$, whereas these principles fail for informational consequence. However, in other ways dynamic consequence diverges even more radically from classical logic than informational consequence. For example, the rule of reiteration, $\phi \vdash \phi$, which is valid in classical logic and almost all well-known nonclassical logics, is not dynamically valid. ‘Might ϕ and not ϕ ’ is not a dynamic consequence of itself, since if s contains both ϕ and ‘Not ϕ ’ worlds, $s[\text{‘Might } \phi \text{ and not } \phi\text{’}] = s[\text{‘Might } \phi\text{’}][\text{‘Not } \phi\text{’}] = s[\text{‘Not } \phi\text{’}]$, a nonempty set of worlds which fails to accept ‘Might ϕ ’ and hence fails to accept ‘Might ϕ and not ϕ ’.

¹³ Sorensen (2009) also notes this order-sensitivity.

- (5) a. If Jack and Jill fail and they might pass, then their preparation is to blame
- b. If Jack and Jill might pass and they fail, then their preparation is to blame

Yalcin's account assigns the same semantic value to (5b) and (5a), and hence fails to explain this contrast. On the other hand, there are other cases where the order of conjuncts seems to make little difference. For example, our reactions to 'If it might not be raining and is raining ...' do not seem markedly different from our reactions to Yalcin's 'If it is raining and might not be raining ...'. The variable character of our reactions looks hard to explain from either standpoint.¹⁴

A further difference between the two approaches emerges when we turn to disjunctions. The contrast illustrated by (3a) and (3b) arises in a very similar way for disjunctions:

- (6) a. # Either it's raining and it might not be, or it's a good day for a picnic
- b. Either it's raining and I don't know it, or it's a good day for a picnic

If we assume a truth-functional semantics for 'or', Yalcin's semantics suggests no particular reason why (6a) should be problematic. An explanation of the badness of (6a) that paralleled Yalcin's account of (3a) would need to rely on some non-truth-functional semantics for 'or', similar in relevant respects to Yalcin's semantics for 'if'. But the obstacles to such a project should not be underestimated. To mention just one, recall that for Yalcin, 'If ϕ, ψ ' is world-insensitive: when it is true relative to $\langle s, w \rangle$, it is true relative to $\langle s, w' \rangle$ for any w' . But it would

¹⁴ 'If it might be raining but in fact it isn't raining ...' and 'If it might be raining but it turns out not to be raining ...' are considerably better than 'If it might be raining and it isn't raining ...'. One could try to explain this by devising special semantic values for 'but', 'in fact', or 'it turns out that', but the project does not look promising.

be very odd to treat disjunctions as world-insensitive. For example, this would assign different semantic profiles to ‘ ϕ or ϕ ’ and ϕ .¹⁵

The dynamic account, by contrast, extends more naturally to (6a). Since the first disjunct is not true relative to any $\langle s, w \rangle$, any plausible semantic account of ‘or’ should entail that it is redundant; this is arguably all we need to account for the infelicity of (6a).

While this point favours the dynamic account, it is worth noting that because the dynamic explanation depends on the order of the conjuncts, it does not carry over to (6c):

(6) c. # Either it might not be raining and is raining, or it’s a good day for a picnic

This is a problem, since there is no clear difference in felicity between (6a) and (6c). Thus neither of the accounts on the table so far suggests any account of the infelicity of (6c).

3.2 *Shifted uses*

There are cases where it is tempting to treat ‘It might be that ϕ ’ as equivalent to ‘ α does not know that not ϕ ’ (or ‘ α is not in a position to know that not ϕ ’), where α refers to some subject S who is neither the speaker nor a group including the speaker. Some examples:

¹⁵ Even if you could learn to live with this, perhaps by saying that ‘ ϕ or ϕ ’ is like ‘Must ϕ ’ in being logically equivalent to ϕ in the sense of informational consequence, such an approach risks generating disastrous results for embedded uses of ‘or’. For example, ‘Everything that is either F or F is G ’ will be semantically different from ‘Everything that is F is G ’. One might attempt to deal with such problems by treating ‘or’-sentences as ambiguous between a world-sensitive and a world-insensitive interpretation (cf. Kolodny and MacFarlane 2010, n. 19). But even leaving Grice’s razor aside, those who posit such ambiguity will need to explain why the felicity of (6a) is not salvaged by the world-sensitive reading.

- (7) Sally is hiding at the back of the bus because I might be on it (cf. Egan, Hawthorne and Weatherson 2005, example 16)
- (8) Fred is putting a tick beside the name of each student who might be worth admitting
- (9) I do not know whether the late Antarctic spring might be caused by ozone depletion (Yalcin 2007, p. 1013)

These sentences are problematic from the standpoint of Yalcin’s semantics as we have given it up to now. For example, (8) can be used felicitously even by an ignorant speaker who knows nothing about which students are worth admitting and who can see that Fred is not putting a tick beside *every* name. ‘Might be worth admitting’ in this context seems to mean something like ‘is worth admitting as far as Fred knows’. (9), meanwhile, is inconsistent if we force ‘know’ and ‘might’ always to behave in the way described in section 2. Yalcin notes this problem, and proposes the following solution: the natural (consistent) reading of (9) involves ‘a tacit shift in the information parameter under the scope of “knows”’, so that what matters is the information possessed by ‘relevant experts’. Yalcin is not fully explicit about what he means by ‘tacit shift’. But given that the semantics given in section 2 makes (9) false not merely relative to some particular $\langle s, w \rangle$, but relative to every $\langle s, w \rangle$ (recall that attitude reports are information-insensitive), consistency cannot be saved just by shifting which $\langle s, w \rangle$ is ‘operative’. Instead, the ‘tacit shift’ must work at the level of semantics, so that instead of its normal, information-sensitive semantic value, ‘The late Antarctic spring might be caused by ozone depletion’ on this occasion receives a semantic value which assigns truth to $\langle s, w \rangle$ iff some worlds where the late Antarctic spring is caused by ozone depletion are compatible with what relevant experts know at w . In other words, ‘might ϕ ’ can *optionally* be assigned the

kind of semantic value that the original semantics would assign to ‘Relevant experts believe/know that might ϕ ’.¹⁶

Another class of cases where the semantics of section 2 gives bad results are cases where the verb ‘might’ is interpreted using a future or past perspective. This is natural, even compulsory, in many linguistic environments. For example:

(10) Tomorrow I will read the files, and in the evening I will pay special attention to the people who still might be worth admitting

(11) I am lucky that the treasure was in the first place I dug — it might have been anywhere

Note that such temporal shifts are also common for ‘might’s in the antecedents of conditionals:

(12) If Jones might be worth admitting, we will call her in for a special interview next week

¹⁶ MacFarlane (2011) suggests generating ‘shifted’ semantic values by an operation of ‘free enrichment’, whereby we process certain clauses as if they were preceded by a silent ‘for all so-and-so knows ...’. A problem with this approach is that it fails to explain why we are not free to insert such operators wherever we like: ‘Fred is putting a tick beside the name of each student who is worth admitting’ cannot mean ‘Fred is putting a tick beside the name of each student who is worth admitting for all he knows’. Stephenson (2007) accounts for sentences like (7) using a special ‘shifty’ semantic value for ‘because’, and suggests that other shifted readings should be accounted for by positing a ‘because’ that has undergone ellipsis. This does not strike us as plausible for (8) or (9). But for our purposes, it does not matter whether the proponent of an information parameter generates ‘shifted readings’ by appealing to lexical ambiguity, structural ambiguity, free enrichment, ellipsis, or some other mechanism, so long as the readings are generated somehow.

On the most obvious reading of (12), the antecedent concerns our future perspective: indeed, many of the most natural examples of ‘might’s in the antecedents of conditionals are like (12) in this respect. Yalcin does not discuss this phenomenon of temporal shift, but presumably he would need to treat these examples, like (9), as involving the replacement of the usual information-sensitive semantic value of ‘might ϕ ’ with an information-insensitive alternative. Given how common these ‘temporally shifted’ uses of ‘might’ are, this move will leave the original theory with a much smaller range of application than one might have initially been led to suppose.

Prima facie, allowing ‘shifted’ semantic values for ‘Might ϕ ’ undermines Yalcin’s proposed account of the infelicity of sentences like (3a). A conditional that begins ‘If it is raining and it is compatible with what α knows/believes that it is not raining ...’ is perfectly fine. If this is an option that is available to us in interpreting ‘If it is raining and it might not be raining ...’, one might expect that such conditionals should not in fact strike us as infelicitous. After all, when a sentence admits several readings, and one of these readings is vacuously true or vacuously false, our usual reaction is not to dismiss the sentence as unintelligible, but to focus on other readings which are more likely to be intended.¹⁷ Thus, once shifted semantic values are allowed, Yalcin’s explanations will need to be supplemented with some story which explains why, despite the availability of non-pathological shifted readings for sentences like (3a), our reactions are shaped by their pathological non-shifted readings.¹⁸

¹⁷ This point is especially pressing in view of the fact that the dots in Yalcin’s ‘If it is raining and might not be, ...’ are very naturally filled in either by a consequent about the future, like ‘I will get wet’, or by a consequent that encourages a generic interpretation for the whole conditional, like ‘I get wet’. Perhaps he was implicitly assuming that the dots would be filled by something about the present; nevertheless, even with the above continuations, the conditional is far from great, a fact which Yalcin is not well placed to explain.

¹⁸ Note that Yalcin’s explanation of the intuitive acceptability of (9) (‘I do not know whether the late Antarctic spring might be caused by ozone depletion’) involves the idea that because

One might claim that it is simply the absence of any salient person or group to play the role of the implicit subject S that prevents us from saving (3a) in this way. But this explanation works only if ‘shifted’ readings always have to involve some perspective *other than the speaker’s current perspective*: it is hard to see how the speaker could fail to be ‘salient’ in any explanatorily useful sense. But a view that allows shifted readings only when S is not the speaker or the relevant time is not the present looks implausibly ad hoc. Moreover, there is pressure to allow that S sometimes *can* be the speaker even when the relevant time is the present. Consider

- (13) If there is no spider in the closet and I am having a fatal heart attack because there might be, then I am about to die an absurd death

Plausibly, ‘ ϕ because ψ ’ cannot be true relative to any $\langle s, w \rangle$ unless ‘ ϕ and ψ ’ is.¹⁹ If so, assigning a Yalcin-style semantic value to the ‘might’ clause will make it impossible for the antecedent of (13) to be true relative to any $\langle s, w \rangle$. But (13) is not infelicitous in the same way as (3a). If we allow for shifted readings, it is tempting to think that the ‘might’ in (13) is naturally interpreted as shifted to the speaker’s own perspective: this would fit with the fact that shifted readings seem to arise naturally in other explanatory contexts, such as (7). However, if we explain the acceptability of (13) in this way, we will need to say something

the sentence is pathological on its non-shifted reading, we automatically assume that the intended interpretation is a shifted one (involving ‘relevant experts’). He does not explain why the same mechanism of charitable reinterpretation fails to operate in the case of (3a).

¹⁹ Here we are disagreeing with Stephenson (2007).

else about why we do not save the infelicitous conditionals by reading their ‘might ϕ ’ as ‘for all I know, ϕ ’.²⁰

Of course, even without an explanation of why the ‘shifted’ reading of (3a) is not the salient one, Yalcin can still argue for the superiority of his approach over the old-fashioned approach on which ‘might’ always means ‘for all α knows’ or ‘for all α is in a position to know’, on the grounds that his approach generates *a* pathological reading for (3a), while the old-fashioned approach generates none. But as we will see in section 4, there are ample reasons for propositionalists to posit a range of readings for sentences like (3a) wide enough to include some pathological readings.

4. Constraint and inheritance

As we have seen, Yalcin’s semantics and its variants are logically highly revisionary. The precise nature of the required revisions depend on the account of validity one adopts, but all the options require giving up either modus ponens or *reductio*, and some involve even more radical departures from logical orthodoxy. We think these are heavy costs, to be paid only as a last resort. Even if we set such logical considerations aside, and also set aside foundational worries about expressivism and relativism of the sort discussed in Dreier 2004 and Cappelen and Hawthorne 2009, the specific concerns we have raised about the quality of the explanations provided by Yalcin’s semantics and its variants would be enough to motivate an investigation of alternative approaches. But our main goal is not to raise problems for Yalcin’s account, but to show how Yalcin’s data can be explained within a standard semantic framework, in which the semantic values of sentences (relative to an assignment) are propositions. In the present section, we will make a start on this project, by refining the old-

²⁰ Another case where ‘Might ϕ ’ seems to mean ‘For all I know, ϕ ’: ‘The guy at the door thinks I’m being rude in not opening it, because he doesn’t realise that he might be a plague carrier’.

fashioned account of the range of interpretations ‘might’ and ‘must’ can bear, and considering one way in which the linguistic context of an occurrence of ‘might’ or ‘must’ can help make certain interpretations especially natural.

Philosophers are apt to overplay the significance of the contrast between epistemic and non-epistemic modals. On the usual picture, it is necessary and sufficient for the truth of an epistemic possibility claim that the relevant body of knowledge be consistent with the proposition in question, while the truth conditions of a non-epistemic possibility claim have nothing to do with knowledge. We claim that there are, in addition, a range of ‘constrained’ uses of ‘might’ and ‘must’, where the restriction on the modal combines epistemic and non-epistemic considerations. For ‘Might ϕ ’ to be true at w on a constrained use, ϕ must be true at some world w' which is both compatible with what the relevant subjects know (or are in a position to know) at w , and also matches w in some further, contextually relevant respects.²¹

For example, suppose you draw a coin from a bucket containing some normal coins and some double-headed coins. Without looking at the coin, you say

(14) I’m not sure whether this coin might land Tails

Here ‘This coin might land Tails’ cries out for an interpretation where it is true if the coin is normal and false if it is double-headed. It would be odd to respond to this utterance of (14) with (15):

²¹ We take it that the ‘worlds’ quantified over in this analysis must include metaphysically impossible worlds, given that sentences like ‘Hesperus might not be Phosphorus’ are true in some contexts. However, the range of issues raised by the fact that not all metaphysical necessities are known will play no further role in the present paper.

(15) It's obvious that the coin might land Tails, since for all we know it's a normal coin, and normal coins often land Tails when they are tossed

On the other hand, if after drawing the coin you just utter

(16) This coin might land Tails

there are two things you could be saying. You could be making the flat-footed remark for which (15) would be a reasonable justification, or you could be making a risky guess which would be false if it were a double-headed coin. But either way, in asserting (16) you are conveying your lack of knowledge about how the coin will land. It would be misleading to characterize either use as straightforwardly non-epistemic.²²

Following Hawthorne (2007), we suggest that ordinary uses of 'It might be that ϕ ', 'It is possible that ϕ ', etc. often express constrained epistemic possibility. This explains why such claims are often intuitively false in cases where a purely epistemic reading would predict them to be true. For example, if someone said 'Harry might fall', mistakenly thinking that Harry was very near the edge of a cliff when in fact he was quite far from the edge, it is

²² Others have recognized that 'might' admits constrained uses. For example, DeRose (1998) thinks that 'the particle might veer left' has an interpretation which entails both that the relevant group does not know that the particle will not veer left and that it is not causally determined that the particle will not veer left. His way of generating this truth condition involves keeping the analysis of 'Might ϕ ' as 'there is no relevant way in which the relevant group can come to know that not ϕ ', while being ultra-flexible as regards what counts as a 'relevant way of coming to know', so that 'deduction from the laws and initial conditions' can sometimes count as a relevant way of coming to know anything entailed by the laws and initial conditions, even if no finite being could carry out such a deduction. We have no substantive objection to this way of generating constrained truth-conditions, although we find it easier to employ an accessibility relation which combines epistemic and non-epistemic factors. Our 'constrained readings' are also discussed in Hawthorne 2007, where they are called 'danger-theoretic' readings.

natural to feel that they said something false: this can be explained if the proposition expressed requires there to be a world where Harry falls which is both compatible with the speaker's knowledge and accurate as regards Harry's distance from the cliff. Similarly, it is easy to hear 'This washing machine might be defective' as false when based on the fact that a recall has been issued for a certain make of washing machines to which the demonstrated machine does not belong: this can be explained if the worlds where the machine is defective are required to be accurate as regards its make. It is easy to hear 'Sally is hiding at the back of the bus because I might be on it' as false when the speaker never takes buses: this can be explained if the proposition expressed requires at least one world where the speaker is on the bus to be accurate as regards the speaker's bus-taking habits, as well as compatible with Sally's knowledge.

Constrained uses of 'might' and 'must' thus arise naturally when the extra-linguistic context makes salient a particular respect in which epistemically possible worlds can be accurate or inaccurate. But the *linguistic* context of an occurrence of 'might' or 'must' can also raise a particular respect of matching to salience, and thus help to make a constrained interpretation natural. Consider

(17) Either Bob is at home or he must be working

(17) can be felicitously asserted if one knows that Bob is at home or working, does not know that Bob is at home, and does not know that Bob is working. We propose to explain this by claiming that the 'must' in such uses of (17) is constrained in such a way that 'He must be working' is true iff Bob is working in all epistemically possible worlds *that are accurate as regards whether Bob is at home*. The constraint for the modal in the second disjunct is

‘inherited’ from the first disjunct. On this interpretation, (17) is true iff either Bob is at home, or Bob is working in every epistemically possible world where he is not at home.²³

A similar phenomenon occurs with conjunctions, although testing for it requires considering embedded conjunctions. Consider

(18) Bob is in the pub and might be having fun

which seems to entail or implicate (19):

(19) Bob might be having fun in the pub

This much can be accounted for even if ‘might’ is always purely epistemic: we can say that in uttering (18), the speaker implicates that Bob is in the pub in all epistemically possible worlds while asserting that he is having a good time in one of them. But we need further resources to explain the behaviour of (18) in embedded occurrences, such as

(20) Either Bob went to the pub and might be having fun, or he is with Alice

²³ This phenomenon extends to non-epistemic modals. For example, ‘Either Bob went to see his mother last week or he should have written her a letter’ seems true if Bob did not go to see his mother last week, and had an obligation either to do so or to write her a letter, but no obligation to do either. We can explain this by positing that ‘He should have written her a letter’ contributes a proposition that is true iff Bob writes his mother a letter in all the best worlds (according to the relevant ordering source) which are accurate as regards whether he went to see her last week. The same thing happens in ‘Either he didn’t succeed in robbing the bank or he ought to be in jail’. (That is not to say that modals in disjunctions are *always* constrained. In ‘Either Peter Singer’s arguments are no good or we ought to be giving most of our money to charity’, the ‘ought’ seems to mean the same thing that it would normally mean unembedded.)

Suppose that in fact Bob went to the pub without Alice, and that we know that Bob never has fun at the pub, although for all we know he is having fun elsewhere with Alice. On the purely epistemic reading, (20) is true, since both conjuncts of its first disjunct are true. But (20) seems intuitively false under these circumstances. We propose that the ‘might’ in (20) is naturally interpreted as constrained in such a way that the truth of ‘Bob might be having fun’ requires Bob to be having fun in some epistemically possible world that is accurate as regards whether Bob went to the pub.

The same style of explanation could also be applied to some very well-known facts about the behaviour of modals in the consequents of conditionals. In uttering

(21) If this is a zebra, it must be an animal

one is obviously not ruling out scenarios where the object in question is a zebra that for all anyone knows is a zebra-robot.²⁴ This can be explained by saying that the natural reading of ‘It must be an animal’ in (21) is a constrained reading, equivalent to ‘It is an animal in all epistemically possible worlds that are accurate with respect to whether it is a zebra’.²⁵

There is a long tradition of explaining the facts about sentences like (21) by assigning them non-obvious logical forms. In discussions of modal fallacies, it is often said that such sentences have readings where the modal takes wide scope over the conditional, as in ‘Must(if

²⁴ The phenomenon is of course not specific to epistemic modals. ‘If Jones is carrying an automatic weapon, he ought to be in jail’ is acceptable on a natural reading even if we are sure that Jones ought to be both free and weaponless.

²⁵ If we take indicative conditionals to be material conditionals, (21) is thus true at w iff either the animal is not a zebra at w , or at w we know that it is either a non-zebra or an animal. On a straightforward ‘closest world’ approach, (21) is true at w iff the closest world to w where the animal is a zebra is one at which we know it is either a non-zebra or an animal. (Presumably the relevant notion of closeness will guarantee that this world is epistemically possible at w : see Stalnaker 1975).

this is a zebra, this is an animal)'.²⁶ A more recent view, due to Kratzer (2012, Ch. 4), takes the 'if' clause to combine directly with the 'must' rather than functioning as a separate operator.²⁷ The facts about (21) do not on their own constitute a compelling case for such views, since it is not plausible to assign wide scope to the modals in examples involving disjunctions and conjunctions like (17) and (20),²⁸ and not plausible that 'or' and 'and' in these examples are anything other than the dyadic operators they seem to be. Explanations of (21) that appeal to non-obvious logical forms also have trouble with conditionals like (22) whose consequents contain several modals, unless they posit some syntactically mysterious operation which allows a single pronounced 'if' clause to appear several times at the level of logical form:

(22) If this is a zebra, it must be an animal and it's probably hungry

The constraint-based account of (21), by contrast, generalizes straightforwardly to (22): both the 'must' and the 'probably' are naturally interpreted as constrained in such a way that only worlds which are accurate as regards the truth value of 'This is a zebra' are relevant. While this is not the place for a full comparison of our account with explanations that invoke non-

²⁶ This idea goes back at least to the medieval distinction between *necessitas consequentis* and *necessitas consequentiae*.

²⁷ A view like Kratzer's is also, in effect, adopted in the extensive literature on 'dyadic deontic logic' inspired by Chisholm 1963. A central idea in this literature is that 'If ϕ , it ought be that ψ ' cannot adequately be symbolized either as 'Ought($\phi \rightarrow \psi$)' or as ' $\phi \rightarrow$ (ought ψ)', but should instead be regarded as the result of applying a sui generis binary operator to ϕ and ψ . (See Åqvist 2002 for a survey.)

²⁸ Note that 'Either this is a zebra or it's probably a donkey' does not have any reading on which it is equivalent to 'Probably, this is either a zebra or a donkey'.

obvious logical forms, we take these considerations to constitute a strong *prima facie* advantage of our proposal.²⁹

Our proposal, then, is that that ‘might’s and ‘must’s in contexts of the form

- (23) a. ϕ and it might/must be the case that ψ
- b. Either ϕ or it might/must be the case that ψ
- c. If ϕ , it might/must be the case that ψ

are often naturally read as carrying a constraint inherited from ϕ — that is, as equivalent to ‘In some/all epistemically possible worlds that are accurate with regard to whether ϕ , ψ ’. Under this analysis, (23a–c) are logically equivalent to (24a–c):

- (24) a. ϕ and in some/all epistemically possible worlds in which ϕ , ψ
- b. Either ϕ or in some/all epistemically possible worlds in which not ϕ , ψ
- c. If ϕ , then in some/all epistemically possible worlds in which ϕ , ψ

There is an alternative approach to (23a–c) that generates the same truth-conditions by a different mechanism: namely, by taking ‘Might/must ψ ’ to mean ‘In some/all epistemically possible worlds in which ϕ , ψ ’ in (23) and (23), and ‘In some/all epistemically possible

²⁹ In suggesting that (21) should be assigned its obvious logical form, we are in agreement with Gillies (2010). But our account is propositionalist and logically orthodox, whereas Gillies’ strategy turns on a dynamic semantics similar to the one discussed in Sect. 3.1 above, and thus requires a highly revisionary logic.

worlds in which not ϕ , ψ in (23).³⁰ While this approach assigns different truth-conditions to the embedded modal claims in (23a–c), the difference is washed out at the level of the whole sentence. While the difference between this approach and our account based on hereditary constraint is not crucial to our defence of propositionalism, we do think that our way of generating the relevant truth conditions is superior, for several reasons.³¹

First, such an account is *prima facie* less unified and explanatory, since it treats conjunctions and disjunctions differently.

Second, the proposed readings of ‘Must ψ ’ do not entail ψ , and the proposed readings of ‘Might ψ ’ are not entailed by ψ . It seems odd that these intuitively valid entailments should be violated in this way, even if the violations always occur in embedded environments where they do not show up in untoward truth conditions.

Third, the constraint-based proposal fits more naturally with the standard account of the conversational implicatures of disjunctions. Typically, when one asserts that P or Q, one’s utterance carries both a ‘primary’ implicature, that one does not know whether P and does not know whether Q, and a ‘secondary’ implicature, that it is not the case that P and Q.³² Our proposal respects both of these implicatures. Consider a typical utterance of (17) (‘Either Bob is at home or he must be working’). On our account, the second disjunct says that Bob is working in every epistemically possible world that is accurate as regards whether Bob is at home. To know this proposition without knowing whether Bob was at home, one would have

³⁰ The phenomenon illustrated by (23a) is discussed in the linguistics literature under the heading of ‘modal subordination’ (Roberts 1989): a similar phenomenon also occurs in sequences of sentences, and with sequences of modals. The standard approach in that literature is the one that we are arguing against in the following paragraphs.

³¹ The dynamic semantic theories discussed in Sect. 3.1 could be seen as an attempt to deliver truth-conditions like (24a–c) compositionally, by assigning special semantic values to the connectives. This is not quite right, since it misses the dynamic semanticists’ insistence on a non-propositionalist account of the semantic values of sentences. Nevertheless, at least the third, fourth, and fifth of the criticisms below also apply to the dynamic semantic treatment of (23a–c).

³² This terminology is due to Sauerland (2004).

to know that the only epistemically possible worlds are worlds where Bob is working. To know its negation without knowing whether Bob was at home, one would need to know that there are epistemically possible worlds where Bob is at home and not working. So if the speaker does not know whether Bob is at home and does not know whether Bob is working, but does know that Bob is either at home or working, the standard primary implicature of the disjunction is true. And indeed, (17) does strongly suggest that this is the speaker's epistemic situation. Meanwhile, the standard *secondary* implicature of (17) will true on our analysis so long as it is not the case both that Bob is at home and that every epistemically possible world where he is at home is one where he is working. (17) does seem to carry not just this implicature, but the stronger implicature that there is an epistemically possible world where Bob is at home but not working: if one could rule out this possibility, (17) would be a strange and misleading thing to assert. By contrast, if we took 'He must be working' in (17) to express the proposition that Bob is working in every epistemically possible world where he is not at home, then the standard primary implicature will be false on the plausible assumption that the speaker knows whether there are epistemically possible worlds where Bob is neither at home nor working. And if Bob is in fact at home — a case which the speaker certainly does not rule out — the standard secondary implicature will also be false, since both disjuncts are true.

Fourth, the constraint-based proposal accounts for the fact that after an utterance of (23a–c), subsequent unembedded modals can easily pick up on the relevant constrained interpretation. For example, if one follows up (17) with

(25) On reflection, it seems unlikely that Bob is at home. So I'm pretty confident that he must be working.

it looks like ‘He must be working’ is supposed to be interpreted in the same way that it was interpreted in (17). Similarly, if following an utterance of (21), one utters

(26) It’s starting to look like this is a zebra. So it’s starting to look like it must be an animal.

the ‘must’ in ‘It must be an animal’ is certainly constrained; and moreover, this use of ‘must’ seems to pick up directly on its use in the consequent of (21). Our proposal handles this smoothly: what ‘He must be working’ contributes in (17) (namely, that Bob is working in all epistemically possible worlds that are accurate as regards whether he is at home) is indeed something of which one is pretty confident, and no more than pretty confident, if one is pretty confident that Bob is not at home. By contrast, on the opposing view, what ‘He must be working’ contributes in (17) is the proposition that Bob is working in all epistemically possible worlds where he is not at home. This is surely not what one is claiming to be pretty confident of in asserting (25).

Fifth, consider

(27) Either Bob is in the pub and so must be having fun, or he decided to stay at home

The ‘so’ here indicates some sort of explanatory relation. It is easy to see how Bob’s being in the pub could be part of a good explanation of the fact that he is having fun in all epistemically possible worlds which are accurate with respect to whether he is in the pub. By contrast, it is hard to see how his being in the pub could help to explain the fact that he is having fun in all epistemically possible worlds where he is in the pub.

We are not proposing that the *only* admissible interpretations of (23a–c) are hereditarily constrained interpretations. In some settings, other readings — purely epistemic readings, or constrained readings which do not require accuracy as regards the truth value of ϕ — are more salient. For example, when ϕ itself contains a modal, the relevant constraint will often be derived from the prejacent of that modal rather than from ϕ itself. On the most obvious reading for ‘You should do the Atkins diet and you might lose a lot of weight’, the constraint for the second conjunct is provided by ‘You do the Atkins diet’, not ‘You should do the Atkins diet’. This phenomenon is hard to make sense of on any view that explains the typical truth-conditions of (23a–c) using mechanical rules, but unsurprising from our point of view.³³

So far, we have been discussing sentences in which modals occur as second conjuncts or disjuncts, or in the consequents of conditionals. There are also cases where modals are interpreted as carrying a constraint derived from *later* material in the same sentence. For example, if we swap the conjuncts in (20) to get (28), its intuitive truth-conditions do not seem to be affected:

(28) Either Bob might be having fun and went to the pub, or he is with Alice

The natural reading is still the constrained one on which it is false if Bob is not with Alice and the speaker knows he never has fun in the pub.³⁴ However, it does seem, in general, to be harder for inheritance to operate from right to left. This is not hard to explain. Other things

³³ Klinedinst and Rothschild (2012) note this effect, and deal with it by positing a special-purpose exception to the usual dynamic rule for calculating the semantic value of a conjunction. We find this implausible: intuitively, ‘and’ in these sentences means exactly what it normally means.

³⁴ We can also get this effect in disjunctions, such as ‘He must know by now or he is more of an idiot than I took him for’. However, in some other examples of this kind (though not this one!), one is tempted to treat the first disjunct in ‘Either must ϕ or ψ ’ as unconstrained and the second disjunct as a ‘take-back’.

equal, our interpretation of a modal will be influenced by the facts about what is salient when it is uttered. While we can of course go back and rethink the interpretation of earlier material in the light of later developments, we will expect co-operative speakers to prefer modes of expression that do not require this.

This fact already helps to explain why there should be a bias *towards* hereditarily constrained readings for modals in second conjuncts and disjuncts. For a co-operative speaker who intends some alternative to an hereditarily constrained reading for a modal in a conjunction or disjunction should, *ceteris paribus*, choose an order in which the conjunct or disjunct containing the modal comes first, so as to diminish the risk that hearers will mistakenly take the intended reading to be the hereditarily constrained one, or be distracted by the need to consider such readings. And since we expect one another to be co-operative, hearers will tend to take the fact that a speaker put a modal in a second conjunct or disjunct as defeasible evidence that the intended reading was hereditarily constrained.³⁵

Inheritance will be one of our main resources for explaining the oddity of the embedded occurrences of ‘ ϕ and might not ϕ ’ in Yalcin’s examples. For ‘ ϕ and might not ϕ ’ is *inconsistent* if the ‘might’ is hereditarily constrained: on this reading, ‘Might not ϕ ’ is true at w only if ‘Not ϕ ’ is true at some world that matches w as regards the truth value of ϕ , which can only happen when ϕ is false at w . This reading will thus render ‘If ϕ and might not ϕ , ψ ’ pathological in whatever way conditionals with logically impossible antecedents are pathological.

More needs to be said, however. For we have not claimed—and it would not be defensible to claim—that ‘might’s in second conjuncts *must* be hereditarily constrained. So we need to explain why our reactions to ‘If ϕ and might not ϕ , ψ ’ are dominated by the

³⁵ We expect that this sort of explanation can also be applied to conditionals, but the details will depend on one’s theory of their truth-conditions.

pathological truth-conditions generated by inheritance, rather than the non-pathological truth-conditions yielded by other interpretations. One would normally expect interpretative charity to cause us to pass over pathological interpretations, even if they belong to a kind that is often salient and natural.

In the following sections we will make some suggestions about why the bias towards hereditarily constrained interpretations should be strong enough to dominate our reactions to ‘If ϕ and might not ϕ , ψ ’ in spite of the existence of other, non-pathological interpretations. But before embarking on this project, we want to stress that our goal of undermining Yalcin’s case against propositionalism does not depend on the success of our suggestions. For as we pointed out in section 3, if Yalcin accepts that epistemic modals can receive ‘shifted’ interpretations — as he must, to avoid what would otherwise be obvious counterexamples — he too must explain why charity does not lead us to prefer shifted interpretations when the unshifted interpretations are pathological. You might think he could satisfy this demand just by positing, as a primitive semantic rule, that unshifted interpretations enjoy a certain ‘default’ status. Arguably, semantics sometimes does need to postulate this sort of structure within the set of admissible interpretations of an expression.³⁶ But to the extent that this is an acceptable move for Yalcin, it should be just as acceptable for us to stipulate, without further explanation, that the ‘default’ interpretations of embedded modals are hereditarily constrained.

However, we are not ready to give up on the search for an explanation of the bias towards hereditarily constrained readings. We will attempt to provide such an explanation in sections 6 and 7. This explanation will turn on two observations which section 5 will

³⁶ We are thinking, for example, of the fact that although ‘strong’ can express a wide range of properties, many of which have nothing to do with physical strength (e.g. in ‘strong student’), interpretations involving physical strength have a certain default status. This seems to be a lexical matter: ‘powerful’ admits roughly the same interpretations, but with less of a bias towards the physical.

present — one about our tendency to assume that people know about their own knowledge, and one about the implicatures carried by conditionals with conjunctive antecedents and disjunctions with conjunctive disjuncts.

5. Self-knowledge and implicature

Knowledge is not transparent: people sometimes fail to know whether they know. There are compelling arguments that one can know that P without knowing that one knows that P (see Williamson 2000). And there are even more straightforward considerations that show that one can fail to know that P without knowing that one fails to know that P. Nevertheless, in ordinary settings, there is something quite odd about *avowing* that one fails to know whether one knows. Questions like ‘Do you know whether there’s another train this evening?’ expect a yes or no answer. The answer ‘I don’t know’ is naturally heard as equivalent to ‘no’, rather than to the exotic ‘I don’t know whether I know’.

One might try explaining these facts by appealing to the hypothesis that it is impossible to *know* that one does not know whether one knows that P, together with the idea that asserting involves representing that one knows. This hypothesis is not refuted by the standard counterexamples to the claim that one always knows whether one knows that P. In familiar cases where one fails to know without knowing that one fails to know, one wrongly thinks one knows one knows, and is thus unaware of one’s failure to know whether one knows; and Williamson’s examples of knowledge without knowledge of knowledge are most readily imagined as cases where one is unaware of one’s lack of second-order knowledge. But whether the hypothesis is true or not, what is important for our purposes is that we do not normally expect people to assert or implicate that they do not know whether they know something. Making sense of such claims requires self-conscious epistemological reflection;

and we expect speakers who want to provoke such reflection to do a certain amount of preliminary scene-setting.³⁷

Similar remarks apply to many other mental states. So, for example, while there are compelling arguments against the transparency of the state of *being sure*, in conversation we do not expect people to assert or implicate that they do not know whether they are sure that P, or that they are not sure whether they are sure that P.

Given that we do not expect people to assert or implicate that they do not know how much they know, it is unsurprising that the modals in sentences like (29a–c) are not naturally interpreted as purely epistemic and egocentric:

- (29) a. I don't know whether this coin might/must land Heads
b. I am not sure whether this coin might/must land Heads
c. I am moderately confident that this coin might/must land Heads

Hearing such sentences, we naturally look around for some interpretation of the modals as constrained, or as involving an epistemic perspective other than the speaker's, or both. For

³⁷ We are not suggesting that 'I don't know whether I know' is tolerable only in *philosophical* settings. We can also get into a mood where we treat our memories like databases, whose contents we count as knowing, but not as knowing that we know until we have engaged in a search: 'I'm not sure whether I know the capital of Gambia — I haven't had time to think about it'. It is also possible, even outside philosophy, to get into a mood where we count ourselves as ignorant about the precise requirements for a true belief to count as knowledge. ('Are you saying you know he's coming?' 'Well, I don't know whether I *know*, but the evidence is very strong'.) But these are both quite distinctive settings, which require special cues to help the audience get into the relevant frame of mind.

example, if the question whether the coin is normal or double-tailed is salient, it is natural to favour a constrained interpretation that requires accuracy with respect to this question.³⁸

We have already mentioned the familiar ‘primary’ implicatures of disjunctions: ‘Either ϕ or ψ ’ typically implicates that one does not know whether ϕ is true and does not know whether ψ is true. Similarly, ‘If ϕ , ψ ’ typically implicates that one does not know whether ϕ is true. Given these facts, it is no surprise that ‘If might ϕ , ψ ’, ‘If must ϕ , ψ ’, ‘Either might ϕ or ψ ’ and ‘Either must ϕ or ψ ’ tend, like (29a–c), to send us looking for alternatives to an unconstrained egocentric interpretation. For example,

- (30) If these mushrooms might kill us, then they are not the kind of mushrooms you can buy in the supermarket

suggests a constrained interpretation in which the biological facts about the mushrooms are held fixed, so that the antecedent is true iff there are worlds consistent with our knowledge where the mushrooms kill us while having their actual biological characteristics. In other settings, the natural alternatives involve the epistemic perspectives of individuals or groups concerning which the speaker can plausibly be presumed to be ignorant. For example, the ‘might’s in (31) seem to concern the doctor’s knowledge:

- (31) a. If you might be anaemic, the doctor will order a blood test

³⁸ Someone might use the preference for constrained or perspective-shifted interpretations of ‘might’ and ‘must’ in (29a–c) as an argument that transparency is actually *true* for the default interpretations. This view strikes us as unpromising, given the extreme generality of the anti-transparency arguments.

- b. Either you might be anaemic, or this doctor just loves ordering blood tests for no good reason

Conversational implicatures are famously subject to the vagaries of context. The claim that ‘If ϕ , ψ ’ implicates ignorance of the truth value of ϕ is only a *ceteris paribus* generalization. One setting in which such implicatures are blocked involves the use of a conditional to signal that one is reasoning from the saliently known antecedent. Imagine a detective addressing the assembled suspects: ‘As we previously established, the murder took place before noon. If the murder took place before noon, then the butler, who was on the 10 o’clock train from Paddington, could not have been responsible ...’. Here the conditional serves to highlight an inferential connection: given what came before, there is no risk that the audience will take the speaker to be raising some new doubt about whether the murder took place before noon. We will call these uses of conditionals ‘echoing’ uses, although their acceptability does not always require that the antecedent should actually have been asserted.³⁹ Their hallmark is that the usual ignorance implicature is replaced by the implicature that the speaker *does* know the antecedent.⁴⁰

If it were not for the phenomenon of echoing, we would expect ‘If I know that ϕ , ψ ’ and ‘If I don’t know whether ϕ , ψ ’ to carry the unusual implicature that the speaker does not know whether ‘I know that ϕ ’ is true (when the antecedents of the conditionals concern the present time). In fact, however, many such conditionals are naturally heard as echoing conditionals,

³⁹ Echoing conditionals are also available when something has recently made it salient that the antecedent is known. After we both see Sally walk in, you could say ‘If Sally is here, Jed must be here too’. However, the mere fact of ϕ ’s expressing shared background knowledge is not enough to prevent ‘If ϕ , ψ ’ from being heard as raising some new doubt about ϕ .

⁴⁰ Our ‘echoing conditionals’ are what Declerck and Reed (2001) call ‘P-factual conditionals’. They use ‘echoing’ more broadly, to encompass conditionals whose antecedents are recently mooted propositions that the speaker need not endorse.

and as such can be acceptable even when we are operating under the standard presumption of self-knowledge: consider ‘If I know that Angelina is pregnant, she can’t be all that publicity-shy’. This is unsurprising: we would expect co-operative speakers who wanted us to engage with the possibility that they do not know what they know to flag that fact explicitly, so that in the absence of special clues, it will typically be more plausible to assume that the usual ignorance implicatures have been suspended.

In many cases, it is also quite natural to hear ‘If might ϕ , ψ ’ and ‘If must ϕ , ψ ’ as echoing conditionals rather than ignorance-implicating conditionals. (Consider ‘If John might be in the pub right now, we should go in to look for him’ or ‘If x must be less than one, x^2 must be less than x ’.) Can the presumption of self-knowledge explain this too? One might suppose not, on the grounds that the presumption only applies to individual knowledge. There is good reason to think that ‘might’s and ‘must’s are easily interpreted as concerning the knowledge of some group including the speaker, even when no particular such group is especially salient.⁴¹ But there is certainly no general presumption that each member of a group knows whether the group knows whether P . Assuming a group knows whatever any of its members know, any case where you know *you* do not know whether P but do not know whether other members of the group know whether P will be one where you do not know whether the group knows whether P ; and you could easily know that you are in such a case. Indeed, when the possibility that the speaker is ignorant about a relevant group’s knowledge is already salient, we have no trouble hearing ‘If might/must ϕ , ψ ’ as a standard ignorance-implicating conditional. If our platoon has just regrouped after splitting up to search a building, ‘If the enemy might be hiding in the basement, at least four of us had better go back down there’ will implicate that the speaker is unsure how much the platoon knows. But without such contextual pointers, we seem to find it more straightforward to process ‘If might ϕ , ψ ’ and ‘If

⁴¹ See De Rose 1991; von Stechow and Gillies 2011.

must ϕ, ψ ' as echoing conditionals than to engage with the hypothesis that the modal concerns the knowledge of some person or group S such that the speaker does not know whether S knows that ϕ is false, even where S is a group including the speaker. Presumably, situations where one is (knowingly) ignorant as regards the knowledge of a relevant group to which one belongs are sufficiently abnormal that co-operative speakers who want to convey that they are in such a situation have reason to favour some more explicit means than an out-of-the-blue utterance of 'If might/must ϕ, ψ '.

The fact that conditionals implicate ignorance of their antecedents does not yet account for the difficulty of accessing the unconstrained egocentric reading of the 'might' in Yalcin's conditionals, 'If ϕ and might not ϕ, ψ '. So long as I know that I do not know whether it is raining, I will not know whether 'It is raining and I don't know that it is raining' is true; so the implicature that I am ignorant of the truth-value of the antecedent will not block an unconstrained egocentric interpretation of the 'might' in 'If it is raining and might not be raining, ...'. To make more progress, we will need to bring to bear our second observation, which concerns the ignorance implicatures of conditionals with conjunctive antecedents and disjunctions with conjunctive disjuncts.

Such conditionals and disjunctions seem to carry further ignorance implicatures that go beyond the standard ones, as illustrated by (32) and (33):

(32) Either Frank is in Rome and he's with his wife, or he's gone AWOL

(33) If John has drunk three beers already and that's a beer he is drinking right now, then he's going to be drunk soon

Asserting (32) tends to implicate that one does not know either that Frank is in Rome or that he is with his wife. Asserting (33) tends to implicate that one does not know either that John has drunk three beers already or that that is a beer he is drinking.

Why should this be? Grice's explanation of the standard ignorance implicatures of conditionals and disjunctions is a good starting point. For Grice, these arise from the presumption that co-operative speakers will obey the 'maxim of quantity', asserting relevant stronger propositions rather than weaker ones when they can do so without sacrificing brevity. If one knew that P or knew that Q, why did one not assert one of those, rather than going for the pointlessly weak 'P or Q' or 'If P then Q'?⁴² Such explanations extend naturally to the implicatures of (32) and (33). Let us focus on (32). It would be uncooperative to utter (32) if you knew that Frank was in Rome, since you could instead assert the stronger

(34) He is in Rome, and he is either with his wife or gone AWOL

Similarly, if you knew Frank was with his wife, you could have strengthened (32) to (35):

(35) He is with his wife, and he is either in Rome or gone AWOL

(34) and (35) are stronger than (32) without being any more complex or bringing in any new potentially irrelevant considerations. Asserting (32) thus implicates that one does not know enough to assert (34) or (35). Moreover, since the second conjuncts of (34) and (35) are obvious logical consequences of (32), your failure to know (34) and (35) despite knowing

⁴² Prima facie, this explanation depends on Grice's view that 'if' is a material conditional. Nevertheless, it would be surprising if the true explanation of the primary implicatures of conditionals bore no resemblance to Grice's.

(32) requires you not to know the first conjuncts of (34) and (35).⁴³ The implicatures of (33) can be accounted for in exactly the same way.

This kind of application of the maxim of quantity has well-known limitations. Life is full of occasions where it is perfectly fine to assert weak propositions even though one could easily have asserted stronger ones. But there is certainly a kernel of truth to Grice's account of ignorance implicatures. And it is very plausible that any suitably worked-out story about why Grice's account works as well as it does for the simpler sentences will carry over to the more complex ones, which seem very much of a piece.⁴⁴

The generalization that disjunctions with conjunctive disjuncts and conditionals with conjunctive antecedents tend to implicate ignorance concerning both conjuncts is of course defeasible. It does not apply to echoing uses of conditionals, where the usual implicature of ignorance is replaced by an implicature of knowledge. There is also a form of echoing which affects only the first conjunct of a conjunctive antecedent. When someone has just asserted ϕ and we want to draw out its consequences, we can assert 'If ϕ and ψ then χ ' without suggesting that there are new reasons to doubt ϕ : 'We have established that the murder took place before noon. If the murder took place before noon and the butler didn't leave Paddington until 10 a.m., then he can't be the killer. So, let's find out when he bought his ticket.' Here the implicature from 'If ϕ and ψ , χ ' to 'I don't know whether ϕ ' is blocked, although the implicature to 'I don't know whether (ϕ and ψ)' remains in place. Note that there is a contrast here between first and second conjuncts: it is much harder for the partial form of

⁴³ The influential neo-Gricean account in Sauerland 2004 does not readily accommodate this explanation, since in Sauerland's theory (34) and (35) do not belong to the 'scalar alternative set' which provides the relevant competitors for (32). However, Spector (2007) presents a more general and explanatory neo-Gricean story on which (34) and (35) do belong to (32)'s alternative set.

⁴⁴ The central controversy in the current literature on scalar implicature concerns whether competition effects can explain *secondary* implicatures, such as the implicature from ' ϕ or ψ ' to 'not (ϕ and ψ)'. Both sides seem to accept that something broadly like Grice's story works for primary implicatures: see Fox 2006, Sect. 1.

echoing to block the implicature from ‘If ϕ and ψ , χ ’ to ‘I don’t know whether ϕ ’, while leaving in place the implicature to ‘I don’t know whether (ϕ and ψ)’. The following speech sounds much more awkward: ‘We have established that the butler didn’t leave Paddington until 10 a.m. If the murder took place before noon and the butler didn’t catch the train from Paddington until 10 a.m., then he can’t be the killer. So, let’s look for more evidence about the time of the murder.’ The fact that order matters in this way can plausibly be explained by appeal to the Gricean maxim of manner (‘Be orderly’).⁴⁵

6. Explaining the pull of inheritance

We are now in a position to explain why there is such strong pressure to interpret the modals in ‘If ϕ and might ψ , χ ’ and ‘Either ϕ and might ψ , or χ ’ as hereditarily constrained. Because such conditionals and disjunctions will generally implicate that the speaker is ignorant of the truth value of ‘might ψ ’, the presumption of transparency will generate pressure not to interpret ‘might’ as unconstrained and egocentric. The relevant difference between these cases and ‘If might ψ , χ ’/‘Either might ψ or χ ’ is that there is no need for us to go hunting around

⁴⁵ Another way for the typical implicatures of conditionals with conjunctive antecedents and disjunctions with conjunctive disjuncts to be blocked involves presuppositions. Consider

- (*) If Fred embezzled money and doesn’t regret taking other peoples’ hard-earned property, he is a monster

This does not seem to implicate ignorance of the second conjunct: if we were not sure whether Fred had embezzled money, we could acceptably assert (*) even if we knew that he had no regrets at all about his life. This can plausibly be explained by the fact that ‘ x doesn’t regret ϕ ’ presupposes the truth of ϕ , which gives us a strong reason to avoid competitors like

- (**) Since Fred doesn’t regret taking other peoples’ hard-earned property, if he embezzled money, he is a monster

Because (*) does not carry the unwanted presupposition that Fred took the money (it is ‘filtered out’ by the conditional), it prevails in the competition with (**).

for a contextually relevant alternative interpretation which escapes the presumption of transparency. A salient alternative is guaranteed to be ready to hand, namely a hereditarily constrained interpretation on which ‘might ψ ’ is true at w iff some ψ -world is both consistent with the knowledge available to the speaker (or some other relevant person or group) at w , and agrees with w as regards the truth value of ϕ . Even if the speaker knows exactly which worlds are consistent with the relevant subject’s knowledge, she will be ignorant of the truth value of ‘might ψ ’ on this interpretation so long as she is ignorant of the truth value of ϕ , and the set of worlds in question includes some ψ -worlds but does not include both ‘ ϕ and ψ ’ and ‘Not ϕ and ψ ’ worlds.

This explanation of the pull of inheritance predicts that when there is a salient interpretation of the ‘might’ as involving the unknown facts about the knowledge of some person or group, or as constrained in some non-hereditary way, the modals in ‘If ϕ and might ψ , χ ’ need not be interpreted as hereditarily constrained. This seems indeed to be the case. Consider

- (36) If we get stuck in traffic on the way to the airport and your laptop might have a bomb in it, we’re going to miss our flight

It is natural to take this ‘might’ as having to do with the future knowledge of the TSA agents who are in charge of searching peoples’ bags at the airport. Because of this, there is no special pressure towards a hereditarily constrained interpretation. If hereditary constraint applied along with the shift of perspective, (36) would implicate that for all you know, there is a world compatible with the TSA agents’ knowledge where your laptop has a bomb in it, and which is accurate with respect to whether we get stuck in traffic on the way to the airport. But

in fact, the most salient interpretation of (36) is acceptable even if you know that the TSA agents know that terrorists always take the train.

Let us turn, finally, to Yalcin's conditionals, 'If ϕ and might not ϕ , ψ '. In a standard setting, a conditional of this form will implicate that the speaker is ignorant of the truth value of 'Might not ϕ '. The presumption of transparency will thus steer us away from unconstrained egocentric interpretations of the 'might'. There is bound to be a salient alternative to such an interpretation, namely a hereditarily constrained interpretation. But as we noted in section 4, ' ϕ and might not ϕ ' is *inconsistent* on such an interpretation: the second conjunct is true only at worlds where ϕ is false, since only such worlds can match any world where 'Not ϕ ' is true as regards the truth-value of ϕ . This, we suggest, explains why our reaction to the sentence is dominated by the feeling of pathology generated by the hereditarily constrained interpretation, at least when the speaker has done nothing to raise any alternative interpretation to salience. The same explanation applies *mutatis mutandis* to 'Either ϕ and might not ϕ , or ψ '.

As we observed in section 5, 'If might ϕ , ψ ' is felicitous even when no interpretation that escapes the presumption of transparency is salient: we naturally treat such uses as echoing uses. But echoing cannot save the unconstrained egocentric interpretation of 'If ϕ and might not ϕ , ...'. A normal echoing conditional implicates that the speaker knows the antecedent, but it is impossible for a speaker to know the proposition expressed by ' ϕ and might not ϕ ' on the unconstrained egocentric interpretation, since the second conjunct is false whenever the first is known. And as previously noted, the partial form of echoing that targets only one conjunct applies far more readily when the echoed conjunct comes first. Thus there is no easy escape from the implicature that the speaker does not know whether 'might not ϕ ' is true.⁴⁶

⁴⁶ There are some examples where 'If ϕ and might not ϕ , ψ ' seems to be acceptable on an unconstrained egocentric interpretation. Suppose an intruder walks in on A and B, brandishing a realistic-looking gun. A: 'I'm going to shoot this guy!' B: 'Hold on — if you shoot him and that's not a real gun, you're guilty of murder!' A: 'You're wrong: if that's not a

Two elements of this account are sensitive to the difference between ‘If ϕ and might not ϕ , ψ ’ and ‘If might not ϕ and ϕ , ψ ’. First, hereditarily constrained interpretations are more salient when the constraint is provided by material earlier in the sentence. Second, the partial form of echoing can block the ignorance implicature carried by the first conjunct in a conjunctive antecedent, whereas it is hard or impossible for it to target the second conjunct. These two contrasts seem to provide quite a good fit for the order-sensitivity in the data. While ‘If it might not be raining but it is raining, ... ’ is odd out of the blue, we find it does not so strongly elicit the phenomenology of contradiction as (3a) (‘If it is raining but might not be, ... ’). Our reaction is more ‘What on earth are they trying to say?’ than ‘How is that even possible?’. And given an appropriate consequent, it is not so hard to save the order-reversed conditionals by invoking partial echoing:

(37) If it might be raining but it is in fact dry, then you were wrong when you predicted that whenever the weather cleared up we would know it

Imagined as a follow-up to ‘It might be dry and it might be raining’, (37) is not so odd.

One might have thought, offhand, that no diagnosis of the badness of (3a) which depends on facts about implicature could explain why it is so much worse than (3b) (‘If it is raining and I don’t know that it is raining, then ... ’). Given the general pattern we have observed for

real gun but it might be, I’ll only be guilty of manslaughter’. Here ‘it might be’ seems to express the proposition that for all A knows it is a real gun. The usual ignorance implicature seems to be absent — there is little suggestion that A does not know whether the gun might be real. We suggest that in this case, A’s speech is licensed by the desire to highlight a general principle such as ‘If you shoot someone who might be brandishing a real gun, you are at most guilty of manslaughter’. Competitor sentences like ‘Since it might be a real gun, if it’s not a real gun I’ll only be guilty of manslaughter’, which we would normally expect to be preferred on Gricean grounds, are in this case worse because they do not so effectively highlight the relevant general principle.

conditionals with conjunctive antecedents, we would expect (3b)—assuming a completion that makes it clear that the antecedent is about the present, such as ‘... then I am about to lose my reputation as a weather-watcher’—to implicate that the speaker does not know whether she knows that it is raining, and hence to be rather odd or jarring outside certain specialized contexts. But this prediction does not undermine the contrast with (3a). The role of implicature in our account of (3a) is to explain why we do readily see an alternative to the hereditarily constrained interpretation, given the pathological truth conditions it generates. But (3b) has no interpretation where its antecedent is contradictory: at worst, it forces us to enter into an unusually sophisticated mode of reflection in which we take the speaker to think she does not know whether she knows. While this may come as a jolt if we are not already in a philosophical mood, there need be no remaining sense of linguistic oddness to (3b) once we have got into the right mood.⁴⁷

If we left the discussion here it would be seriously incomplete. For—as we saw in discussing sentences like ‘I don’t know whether this coin might land Heads’—hearers are often quite good at figuring out that the intended interpretation of an epistemic modal is some sort of perspective-shifted or constrained interpretation, even when the speaker has provided no helpful cues to draw our attention to any particular such interpretation. While the absence of such cues sometimes makes us hear ‘If might ϕ , ψ ’ as an echoing conditional, on other occasions we conclude that some non-transparent interpretation must be intended and hope that the subsequent course of the discourse will help us understand what it is. One might

⁴⁷ In fact the situation with (3b) is a bit better than that. As discussed in n. 45, the desire to filter out unwanted presuppositions can sometimes block the usual ignorance implicatures of conditionals with conjunctive antecedents. Given the widely shared view that ‘ x doesn’t know that ϕ ’ presupposes the truth of ϕ , n. 45’s account of examples with ‘regret’ carries over to examples like (3b). Indeed, if we replace ‘know’ with ‘believe’ (which carries no presuppositions), the implication of ignorance becomes clearer: ‘If it is raining right now but I don’t believe that it is, I am about to lose my reputation as a weather-watcher’ quite strongly suggests that the speaker does not know what she believes.

expect us to be similarly resourceful when confronted with ‘If ϕ and might not ϕ , ψ ’, charitably setting aside *both* the unconstrained egocentric interpretation and hereditarily constrained interpretations in favour of some yet-to-be-determined alternative. Section 7 will attempt to explain our failure to exercise such creativity.

7. Preferences for explicitness

Let us begin by focusing on cases like those discussed in section 3.2, where there are plenty of clues that the intended interpretation of ‘might’ or ‘must’ involves the knowledge of some third party (or the speaker’s knowledge at some past or future time), so that no presumption of transparency applies. It is important to observe that even in these settings, there is something quite odd about ‘ ϕ and might not ϕ ’ — a fact that cannot be explained by any special semantics for ‘might’ that, like Yalcin’s, only applies to non-shifted uses. For example, (38) and (39) sound odd despite the fact that it is obvious that the relevant epistemic perspective is Sally’s:

(38) Sally is going to waste a lot of time searching for Fred on that bus, and that’s because he isn’t on the bus but might be

(39) Sally is going to waste a lot of time searching for Fred on the bus, if he isn’t on the bus but might be

These are not as crashingly bad as out-of-the-blue utterances of ‘If ϕ and might not ϕ , ψ ’, but they are quite awkward nevertheless, and clearly worse than the corresponding sentences with ‘She thinks he might be on the bus’ or ‘She doesn’t know he isn’t on the bus’. This is puzzling, given that a Sally-relative reading of the ‘might’ is as salient as one could wish.

One might be tempted to explain the oddity of (38) and (39) by appealing to some story on which inheritance-like effects are generated by a mechanical process which operates automatically in all cases. But this seems unpromising: we have already seen several examples where inheritance is optional (e.g. (36) in Sect. 6), and moreover (38) and (39) remain odd when we reverse the order of the clauses separated by ‘but’.

We want to suggest that this residual oddity is part of a wider pattern whereby under certain circumstances, sentences involving expressions that introduce implicit, contextually-supplied arguments are markedly worse than counterpart sentences involving expressions that make the relevant arguments explicit.⁴⁸ Consider the following pairs:

(40) a. It’s not raining over there but it’s raining here

b. ? It’s not raining over there but it’s raining

(41) a. Jim knew there would be good weather somewhere but he was surprised it wasn’t raining where he was

b. ? Jim knew there would be good weather somewhere but he was surprised it wasn’t raining

(42) a. Although Jim’s friends expected him to propose marriage to me, I didn’t expect the proposal

b. ? Although Jim’s friends expected him to propose marriage to me, the proposal was unexpected

⁴⁸ There has been much debate about how ‘implicit reference’ should be implemented in semantic theory. Some argue that implicit reference is always due to unpronounced but syntactically real variables (Stanley 2000), while others invoke more flexible theoretical tools like ‘free enrichment’ (Recanati 2002). The resolution of this debate does not matter for our purposes.

- (43) a. Mary is an enemy of my father's, and worse still, she is an enemy of mine
b. ? Mary is an enemy of my father's, and worse still, she is an enemy

What makes the (b) sentences so strange? We take it that they all involve expressions that when uttered without an explicit argument, can function semantically as if such an argument were present. 'It's raining' can mean 'It's raining here'; 'unexpected' can mean 'unexpected to me', and so on. And considerations of salience cannot by themselves explain why we would have any trouble providing the implicit arguments which would make the (b) sentences semantically equivalent to the (a) sentences — in (42b), for example, the speaker is made completely salient by the occurrence of 'me'.

A natural story about these examples is that the nature of the information being conveyed makes it desirable to place linguistic *focus* on a certain argument-place. The purpose of such focus, in general, is to highlight a range of alternatives to the proposition semantically expressed by a certain sentence, which differ from that proposition as regards the value of the focused constituent.⁴⁹ Since focus is standardly achieved by means of a certain kind of stress, it is not available when the relevant argument is not phonologically realized. In all of the above examples, the point of the two conjuncts is to highlight some difference or similarity in truth value between propositions that differ only in one argument place, so there is strong pressure to find a mode of expression which makes that argument place explicit, thereby allowing it to be focused. (In 'It is not cold here but it is raining', on the other hand, the contrast is not along the parameter introduced by the location argument, and so there is no similar pressure to make that argument explicit.)

'Might' gives rise to similar contrasts:

⁴⁹ This idea about the role of focus is due to Rooth (1992). Glanzberg 2005 is a helpful introduction.

- (44) a. Although you and I know that Fred is not on Sally's bus, he might be on it as far as Sally is concerned
- b. ? Although you and I know that Fred is not on Sally's bus, he might be on it

Even when it is obvious that the relevant perspective is Sally's — for example, considered as an answer to the question 'Why is Sally wasting her time looking for Fred on the bus?' — (44b) is quite odd.

There are many other kinds of cases where explicit reference is strongly preferred for reasons that cannot helpfully be explained in terms of salience. In the following examples, the (b) sentences seem to be infelicitous as ways of communicating the content of the (a) sentences, whereas the (c) sentences are fine:

- (45) a. I was fired from my job as a soup-taster because the week-old borscht tasted so good to me
- b. ? I was fired from my job as a soup-taster because the week-old borscht was so tasty
- c. I didn't bother making new borscht because the week-old borscht was so tasty
- (46) a. Mary will be thought dull if she isn't interested in classic literature
- b. ? Mary will be thought dull if classic literature isn't interesting
- c. Mary will switch to science if classic literature isn't interesting

Why should the explicit (a) sentences be preferred to the implicit (b) sentences, given the felicity of the structurally similar (c) sentences? While there may be several factors at work, we tentatively suggest that considerations of contrast are once again relevant.⁵⁰ Unlike (45c), the explanation in (45a) turns on the fact that the week-old borscht would not have tasted good to others; similarly, to grasp the grounds for the conditional in (46a), the hearer needs to attend to a possible contrast between Mary's attitude to classic literature and those of others.⁵¹ Although there is no need to *focus* 'me' in (45a) or 'she' in (46a), it is not implausible to suppose that hearers will have more trouble latching on to the relevant set of alternatives when the argument place with respect to which the alternatives differ is left unpronounced.

Similar contrasts also arise with epistemic modals. In the following examples, the (b) sentences are quite awkward considered as ways of communicating the content of the (a) sentences, although the structurally similar (c) sentences show that the relevant perspective-shifted readings are available:

- (47) a. Sally needs better glasses if she doesn't know that Fred isn't on that bus
- b. ? Sally needs better glasses if Fred might be on that bus
- c. Sally will wait for the next bus if Fred might be on that bus

⁵⁰ Relativists about 'tasty' will of course claim that 'the week-old borscht was tasty' has an interpretation that is not equivalent to anything of the form 'the week-old borscht tasted good to α '. But it is hard to see how this could help explain our contrasts. There is pressure for relativists to concede that 'tasty' can *sometimes* mean 'tasty for α ', and indeed occurrences of 'tasty' under 'because' are often especially good candidates for this contextualist treatment (Stephenson 2007). So long as there is *any* reading of (45b) that makes it equivalent to (45a), we will need some new resources to explain the oddity of (45b).

⁵¹ Note that the oddity of (45b) depends on the fact that the envisaged causal mechanism involves the employers realizing that their employee does not have the right tastes in soup for the job. It is far better if we are imagining a scenario where the speaker eats all the borscht because it tastes so good to her, and is then fired because tasters are expected to have just a spoonful of each soup.

- (48) a. The TSA agent showed himself to be incompetent because he didn't know that Mrs Smith wasn't a terrorist
- b. ? The TSA agent showed himself to be incompetent because Mrs Smith might have been a terrorist
- c. The TSA agent needed to do a full-body search because Mrs Smith might have been a terrorist

Again, contrast seems to be playing a role: the claims being made by the (a) sentences, unlike those made by the (c) sentences, turn crucially on the idea that different people have different epistemic abilities, which lead to differences in their knowledge even when they are placed in similar external circumstances.⁵²

The idea that pronouncing an argument explicitly helps to activate a certain range of alternatives can also help explain why 'ϕ and might not ϕ' and 'Might not ϕ and ϕ' remain awkward even when the context makes salient an interpretation where the 'might' concerns a third party's knowledge (or a non-present time). Consider (49b), which as we observed at the beginning of this section is quite awkward even when the surrounding discourse makes it quite clear that the intended interpretation is one equivalent to (49a):

- (49) a. Fred is not on that bus and Sally doesn't know it
- b. ? Fred is not on that bus and he might be

⁵² As with (45b) (see n. 51), the oddity of (48b) depends on a particular causal route being salient. (48b) becomes much better if we imagine that the TSA agent sprinted over to Mrs Smith because he could not tell she was not a terrorist, and fell flat on his face, thereby revealing his incompetence.

The fact that it is so natural to substitute ‘but’ for the ‘and’ in (49a) strongly suggests that the conjunctive proposition expressed is one that we process as presenting a contrast. In this case, the contrast is not between Sally and other people, but between the way things are and the way they are known to be by Sally: nevertheless, the need to highlight this contrast motivates explicitly mentioning Sally rather than leaving her as the value of an implicit argument.

Whether or not our suggestions about the role of contrast get to the heart of the matter, examples like (40)–(43) and (45)–(46) show that there is a wide range of logically complex thoughts which cry out for modes of expression involving explicit rather than tacit reference. Against this background, it is not so surprising that ‘ ϕ and might not ϕ ’ should remain somewhat challenging, whether bare or embedded, even when accompanied by abundant evidence that the intended reading concerns the knowledge of some subject other the speaker (or time other than the present).⁵³

We suggest that whatever accounts for these facts can also address the puzzle we left hanging at the end of section 6. The challenge there was to explain why we do not rescue ‘If ϕ and might not ϕ , ψ ’ by charitably assuming that the intended interpretation involves some subject S and time t such that the speaker is unsure whether S knows the proposition expressed by ϕ at t . The mere absence of salient contextual pointers is not an adequate explanation, since it does not hold us back from reaching for such interpretations in the case of ‘If might ϕ , ψ ’. Now we can do better: given that ‘ ϕ and might not ϕ ’ would be a bad way

⁵³ Here are a few examples of the form ‘If ϕ and might not ϕ , ψ ’ where the context forces a non-present reading: ‘If there are no diamonds in the room but there might be, I’m going to spend all day tomorrow searching in vain’; ‘If there is no treasure but there might be, an island gets riddled with pits dug by treasure-hunters’; ‘If there are no mines on the road but there might be, we end up taking a pointless detour’. These seem better than the likes of ‘If there are no diamonds in the room right now and there might be, I am wasting my time on this search’, but still not nearly as good as versions with ‘know’, such as ‘If there are no diamonds in the room and I don’t know it, I’m going to spend all day tomorrow searching in vain’. This is what our account predicts, since the intermediate examples will suffer from the kind of infelicity discussed in this section but not from the kind of infelicity discussed in Sect. 6.

of expressing *any* of its admissible interpretations, it is little wonder that after we notice the pathological character of ‘If ϕ and might ϕ , ψ ’ on the (very salient) hereditarily constrained interpretation, we do not assiduously search for a non-pathological alternative.

8. Embedding under quantifiers

When we invoke conversational implicatures to explain the infelicity of a certain class of sentences, we cannot, in general, expect our explanation to carry over to more complicated sentences that embed those sentences. It is thus good methodology to put section 6’s implicature-based explanation to the test by considering certain more complicated sentences to which natural extensions of Yalcin’s semantics would assign pathological semantic values, but where one would not expect implicatures of ignorance. Universal quantification looks like a good test case. Assuming Yalcin treated ‘Every μ is a ν ’ as equivalent to ‘ $\forall x(\text{if } x \text{ is a } \mu, x \text{ is a } \nu)$ ’, his semantics would predict that ‘Every μ that might be a non- μ is a ν ’ is invariably vacuously true, since ‘ x is a μ that might be a non- μ ’ cannot be accepted by any set of worlds on any assignment. By contrast, the portion of our account of Yalcin’s conditionals which depends on ignorance implicatures will not carry over to the quantified sentences, since ‘Every μ is a ν ’ does not normally implicate ‘There is something of which I don’t know whether it is a μ ’. Our account thus predicts that ‘Every μ that might be a non- μ is a ν ’ will not suffer from the severe kind of badness that afflicts ‘If ϕ and might not ϕ , ψ ’ when no perspective-shifted or non-hereditarily constrained reading of the ‘might’ is salient. We would not, however, predict that it would be completely unproblematic, since the preference for explicitness discussed in section 7, which makes ‘ ϕ and might not ϕ ’ worse than ‘ ϕ and α doesn’t know that ϕ ’ irrespective of embedding, will also make ‘Every μ that might be a non- μ is a ν ’ worse than ‘Every μ that isn’t known by α to be a μ is a ν ’.

What level of contrast do we actually find when we consider such pairs?

- (50) a. Everyone who is HIV-negative but might be HIV-positive has to stay in hospital until their negative test results come through
- b. Everyone who is HIV-negative but isn't known to be HIV-negative has to stay in hospital until their negative test results come through

There is no doubt that (50a) is worse than (50b).⁵⁴ But we do not think that the fine-grained facts are very encouraging for a Yalcin-style explanation. For the degree of improvement between (50a) and (50b) seems roughly the same whether we imagine the sentences being uttered by the hospital authorities themselves or by some outsider reporting on the hospital's policies. There is a big difference here between (50a–b) and corresponding conditionals:

- (51) a. If Fred is HIV-negative but might be HIV-positive, he has to stay in hospital until his negative test results come through
- b. If Fred is HIV-negative but not known to be HIV-negative, he has to stay in hospital until his negative test results come through

If we imagine (51a–b) as spoken by an onlooker commenting on the hospital's policies, (51a) is only moderately worse than (51b). The reaction is 'I see what you're trying to say, but that's a strange way to put it'. On the other hand, if we imagine (51a) being uttered by the

⁵⁴ We note that the effects of ordering seem quite pronounced in this case: 'Everyone who might be HIV-positive but is actually HIV-negative will be relieved when they find out' seems pretty good.

doctor, it sounds awful. This is evidence in favour of some account with the same overall shape as ours: one factor that makes arbitrary embeddings of ‘ ϕ and might not ϕ ’ moderately bad whether or not there is a salient perspective-shifted or non-hereditarily constrained interpretation, and a second factor that makes certain particular embeddings (including disjunctions and conditionals, but not quantifiers) very bad, but only in cases where no perspective-shifted or non-hereditarily constrained interpretation is salient.⁵⁵

9. Supposition and other attitudes

We have yet to say anything about Yalcin’s contrast with ‘suppose’:

- (2) a. # Suppose that it is raining and it might not be raining
- b. Suppose that it is raining and I don’t know that it is raining

Like many others, we find it plausible that there is a close semantic connection between the discourse ‘Suppose ϕ . Then ψ .’ and the single sentence ‘If ϕ , ψ ’. We are tempted by an account of ‘Suppose ϕ ’ built around this connection, according to which its function is to supply an implicit antecedent for subsequent occurrences of ‘then’ until further notice: after ‘Suppose ϕ ’, ‘Then ψ ’ will be interpreted as ‘If ϕ then ψ ’.⁵⁶ To the extent that this is a good

⁵⁵ Another embedding environment to consider is that of the epistemic ‘might’ in counterfactuals. Consider ‘If Fred had been HIV-negative but might have been HIV-positive, he would have had a pointless stay in the hospital’. While this is not perfect, we find it markedly better than the corresponding indicative conditional. Just as in the case of quantified sentences, this is what our account predicts: the counterfactual is bad because of the preference for explicitness, but not terrible because implicature-related problems do not arise.

⁵⁶ Although most philosophers who emphasize the relation between ‘suppose’ and ‘if’ (e.g. Barnett 2006) go on to develop some form of ‘no truth value’ theory of conditionals, the

theory of ‘suppose’, there is no need for us to say anything more about (2a–b). If the point of (2a) is to cause subsequent utterances of ‘Then ψ ’ to be processed like ‘If it is raining and might not be raining, ψ ’, our story about the infelicity of such conditionals will carry over automatically to (2a).⁵⁷

This account is compatible with the idea that ‘suppose’ expresses a propositional attitude relation, which ‘Suppose ϕ ’ instructs hearers to bear to the proposition expressed by ϕ . For one could take the relevant attitude to be constituted by a willingness to process subsequent discourse in the manner described above. Interestingly, however, there is syntactic evidence that ‘Suppose ϕ ’ is not merely an imperative version of ‘ x supposes ϕ ’:

- (52) a. Suppose that Gore were president. (How would things be different? ...)
b. *[I suppose/he supposes] that Gore were president
c. *[I am/he is] supposing that Gore were president

The syntactic well-formedness of (52a) is an aberration. If (52a) is just an imperative to be in a certain mental state, why can one not report one’s own or someone else’s obedience to the command using (52b) or (52c)? The ill-formedness of (52b–c) makes it tempting to treat (52a) as a semantically *sui generis*, discourse-modifying cousin of ‘If Gore were president’.

However, leaving aside the peculiarities of ‘suppose’, it is undoubtedly true that for a wide range of attitude verbs χ , ‘ α χ s that (ϕ and might not ϕ)’ sounds noticeably worse than ‘ α χ s that (ϕ and x doesn’t know ϕ)’ or ‘ α χ s that (ϕ and I don’t know ϕ)’. How can

point is compatible with the view that conditionals express propositions just like any other sentence.

⁵⁷ We note, however, that ‘suppose’ does not seem to license negative-polarity items the way ‘if’ does: ‘If she ever finishes that book ...’ is far more felicitous than ‘Suppose she ever finishes that book’.

propositionalists account for this? Because of the preference for explicitness discussed in section 7, we will expect any sentence that embeds a constituent of the form ‘ ϕ and might not ϕ ’ to be somewhat bad. However, in many cases ‘ α χ s that (ϕ and might not ϕ)’ seems *extremely* bad. Our goal in this section is to explain why this is the case for three large categories of attitude verbs.

Our account will assume a straightforward treatment of attitude ascriptions according to which ‘ α χ s that ϕ ’ expresses the proposition that the referent of α stands in the relation expressed by χ to the proposition expressed by ϕ . We propose that by default, epistemic modals under attitude verbs receive a *subject-centric* interpretation — that is, one where the relevant epistemic perspective is that of the subject of the attitude verb. ‘ α χ s that might ϕ ’ will thus normally mean that the person referred to be α bears the relation expressed by χ to a proposition that is true iff he or she does not know (or is not in a position to know, or lacks evidence that entails ...) the negation of the proposition expressed by ϕ .⁵⁸

It would be good to be able to explain *why* the subject-centric interpretation should have this default status. There is no puzzle about why this should be *one* very salient interpretative possibility. But a bare appeal to salience does not explain why subject-centric interpretations are normally more natural in these environments than *speaker-centric* interpretations — speakers are, after all, normally quite salient. We will not attempt to give a fully worked-out explanation. What we do want to point out is that this fact is, like the contrast-related facts surveyed in section 7, an instance of a quite general phenomenon involving expressions with

⁵⁸ Since we think that propositions can be distinct though necessarily equivalent, we are not forced to identify the propositions expressed by ‘might’ sentences with any propositions expressed by ‘know’ sentences. In a structured proposition framework, for example, one might claim that although the propositions expressed by ‘might’ sentences contain as a constituent a certain modal base that is determined by certain knowledge facts, they do not contain the knowledge relation itself as a constituent. Given this, the contextualist has resources for answering the complaint (Yalcin 2011, p. 308) that contextualist semantics wrongly entails that ‘ α thinks it might be raining’ ascribes a ‘second-order’ thought about the subject’s epistemic state.

‘implicit arguments’. When such expressions occur in attitude ascriptions, there is normally a strong presumption that the implicit argument place is filled by the subject of the attitude verb (if it is an argument place for a person). This interpretative default is indicated in (53)–(58):

- (53) Mary thinks that John is a friend (of hers)
- (54) Fred hopes that there are some decent bars that aren’t too far away (from him)
- (55) Jennifer wants to make sure there aren’t any unexpected (by her) delays
- (56) Sylvester expects the concert to be exhausting (to him)
- (57) Sarah reported that the pattern became visible (to her) after five seconds
- (58) Cameron thinks that Marmite is tasty (to him)⁵⁹

The prominence of subject-centric interpretations in these cases is not really very surprising. People think a lot about themselves; and their thoughts about themselves play a pervasive role in the explanation of other facts about them. It is thus to be expected that a sentence that can be interpreted as reporting either a self-directed thought or an other-directed thought will most readily be interpreted in the former way, without special cues. When we want to discuss someone’s thoughts about some other person, we will normally want to make it obvious that we are doing so by using a formulation that mentions the other person explicitly.

In all the above cases non-subject-centric readings (including speaker-centric readings) are possible, with help from the surrounding linguistic and non-linguistic context. Non-subject-centric readings are also possible for epistemic modals under attitude verbs, and in some contexts they are more natural than the default subject-centric readings:

⁵⁹ Relativists about ‘tasty’ will of course dispute our inclusion of (58) on this list. But few will be tempted by relativism for ‘friend’, ‘far away’, etc.

(59) Mary thinks I am putting a tick beside the name of each student who might be worth admitting

(60) Mary expects to be called for questioning, because she has found out that she might be the murderer

As these examples suggest, the contextual factors that make it natural to interpret a ‘might’ in an attitude report from the perspective of someone other than the subject are quite similar to those that make it natural to interpret an unembedded ‘might’ from the perspective of someone other than the speaker.

The facts about the availability of group interpretations also seem quite similar. Just as unembedded ‘might’s can very easily be indexed to groups containing the speaker, ‘might’s in attitude reports can very easily be indexed to groups containing the subject. This feature is also shared with other expressions with implicit arguments: group interpretations are also easy to access for (53)–(58). In each of these cases, the group interpretations that are easily accessed without special cues are those where the group in question is presumed to be *homogeneous* in the relevant respects—they have the same friends, are exhausted by the same things, find the same things tasty, etc. Similarly, the easy cases for group interpretations of ‘might’ involve groups that are presumed to share the relevant knowledge, while cases where the subject is unsure about the knowledge of other members of the relevant group require more contextual priming.

Let us see how, given the preference towards subject-centric interpretation, we can explain the extreme badness of ‘ α χ s that (ϕ and might not ϕ)’ for several large classes of attitude verbs χ . The most straightforward case comprises attitudes which entail or presuppose

or implicate knowledge. On a subject-centric interpretation, ‘ α knows that (ϕ and might not ϕ)’ is inconsistent: because knowledge distributes over conjunction, it entails ‘ α knows that ϕ and α knows that might not ϕ ’; because knowledge is factive, this entails ‘ α knows that ϕ and might not ϕ ’, which is inconsistent on a subject-centric interpretation of the ‘might’.

Given that subject-centric interpretation is only a default rather than an iron rule, it might seem puzzling that we find out-of-the-blue utterances of ‘ α knows that (ϕ and might not ϕ)’ infelicitous, rather than automatically reaching for some as yet unspecified non-subject-centric interpretation, as we rather easily do in other cases, such as ‘ α doesn’t know whether might ϕ ’. This lack of creativity can be explained in the same way as our lack of creativity in the case of ‘If ϕ and might not ϕ , ψ ’, by reference to the preference for explicitness which makes ‘ ϕ and might not ϕ ’ generally worse than ‘ ϕ and α doesn’t know that ϕ ’.

The second category of attitude verbs we will consider comprises those which, while not entailing or implicating knowledge, do suggest that the subject has some other attitude to the corresponding knowledge claim. For example, ‘ α claims that ϕ ’ plausibly entails or implicates something like ‘ α purports to know that ϕ ’. The contradictory nature of ‘ α knows that (ϕ and might not ϕ)’ on the subject-centric interpretation will thus make for a feeling of pathology in ‘ α claims that (ϕ and might not ϕ)’. Of course it is possible for people to purport to have contradictory properties, so the pathology in question is not outright contradictoriness. But this is an abnormal thing to do; if one wanted to report that someone had done it, one would be expected to choose a form of words that leaves less work for the hearer.

The third category of attitude verbs we will discuss comprises those for which ‘ α χ s that ϕ ’ entails or implicates or presupposes that ‘ α knows that ϕ ’ is *false*. This group includes ‘hope’, ‘fear’, ‘imagine’, ‘suspect’, ‘want’, ‘be pretty confident’, and many others. For these verbs, we can give an explanation of the infelicity of ‘ α χ s that (ϕ and might not ϕ)’ that closely parallels section 6’s explanation of the infelicity of ‘If ϕ and might not ϕ , ψ ’. The

basic idea is that these sentences are pathological on hereditarily constrained interpretations of the ‘might’ (since on this reading they ascribe attitudes towards inconsistent propositions); and such interpretations are hard to escape because their most obvious competitors, namely *unconstrained* subject-centric interpretations, generate implicatures which violate our standard assumptions of transparency.

To explain how these implicatures would arise, we need two further plausible generalizations about this group of attitude verbs. The first generalization is that when they take conjunctive complements, the implicated lack of knowledge extends to each conjunct: ‘ α χ s that (ϕ and ψ)’ implicates not only ‘ α doesn’t know that (ϕ and ψ)’, but also ‘ α doesn’t know that ϕ ’ and ‘ α doesn’t know that ψ ’. For example, if you think James knows that Frank is in the park, it would be quite misleading to say ‘James hopes/fears/suspects/is pretty confident that Mary is in the pub and Frank is in the park’: the conjunctive report strongly suggests some kind of parity in the subject’s epistemic position with regard to the two conjuncts. The second generalization is that the implicatures generated are not only that the subject *in fact* fails to know certain things, but also to the effect that they are (at least implicitly) *aware* of their failure to know these things. Thus if James merely *thinks*, wrongly, that he knows Frank is in the park, ‘James hopes that Frank is in the park’ is still misleading, and so is ‘James hopes that Mary is in the pub and Frank is in the park’. (This generalization plausibly flows from the default presumption of transparency discussed in Sect. 5.)

Given the first generalization, ‘ α χ s that (ϕ and might not ϕ)’ will implicate both (a) ‘ α doesn’t know that ϕ ’ and (b) ‘ α doesn’t know that might not ϕ ’. Given the second generalization, implicature (a) can be strengthened to ‘ α knows that α doesn’t know that ϕ ’. But on a subject-centric reading of the ‘might’, implicature (b) is tantamount to ‘ α doesn’t know that α doesn’t know that ϕ ’, which contradicts the strengthened version of implicature (a). The upshot is that when χ is an attitude verb in the third category, the interaction of

implicature with assumptions about transparency will send us looking for alternatives to a subject-centric interpretation of the ‘might’ in ‘ α χ s that (ϕ and might not ϕ)’. When no helpful contextual cues have been provided, the only salient alternative will be a hereditarily constrained interpretation, under which ‘ ϕ and might not ϕ ’ is itself contradictory.

We will not attempt to argue that all attitude verbs fall into one of the three categories we have considered.⁶⁰ But we have not been able to think of any clear cases where ‘ α χ s that (ϕ and might not ϕ)’ is worse than the considerations discussed above would lead one to expect.

10. Conclusion

Our account of the facts about the felicity of epistemic modals embedded in different linguistic environments has drawn on explanatory resources of several different kinds. We will not pretend to be confident that our way of dividing the work between these resources has struck exactly the right balance. For example, we are open to being convinced that we should have placed more, or less, emphasis on the ‘preference for explicitness’. But since the facts to be explained involve different gradations of infelicity, we are not embarrassed by the multi-layered character of our account of them. The facts strike us as too subtle to be helpfully explained by the mechanical working out of simple compositional rules, even acting on semantic values of a novel sort. Of course, proponents of revisionary approaches to semantics are free to help themselves to many of the explanatory resources we have

⁶⁰ Some verbs straddle the second and third categories. For example, some uses of ‘ α believes that ϕ ’ implicate ‘ α believes that α knows that ϕ ’, while others—especially those with contrastive focus on ‘believes’—implicate ‘ α does not know that ϕ ’. (The latter uses also seem to implicate ‘ α knows that α does not know that ϕ ’, in accordance with our second generalization.) ‘Assume’ and ‘suppose’ are like ‘believe’ in this regard.

developed. But we hope our discussion has shown that there is no pressing need to depart from a logically and semantically conservative approach to epistemic modals.⁶¹

References

- Aloni, Maria, Alastair Butler, and Paul Dekker (eds) 2007: *Questions in Dynamic Semantics*. Kidlington: Elsevier.
- Aloni, Maria, Vadim Kimmelman, Floris Roelofsen, Galit W. Sassoon, Katrin Schulz, and Matthijs Westera (eds) 2011: *Logic, Language and Meaning: Proceedings of the 18th Amsterdam Colloquium. Lecture Notes in Computer Science*, vol. 7218. Heidelberg: Springer.
- Åqvist, Lennart 1984: 'Deontic Logic'. In Gabbay and Guenther 2002, pp. 147–264.
- Barnett, David 2006: 'Zif is IF'. *Mind*, 115, pp. 519–65.
- Cappelen, Herman and John Hawthorne 2009: *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Chisholm, Roderick M. 1963: 'Contrary-to-Duty Imperatives and Deontic Logic'. *Analysis*, 24, pp. 33–6.
- Declerck, Renaat and Susan Reed 2001: *Conditionals: A Comprehensive Empirical Analysis*. Berlin: Mouton de Gruyter.
- DeRose, Keith 1991: 'Epistemic Possibility'. *Philosophical Review*, 100, pp. 581–605.
- 1998: 'Simple "Might"s, Indicative Possibilities and the Open Future'. *Philosophical Quarterly*, 48, pp. 67–82.
- Dreier, Jamie 2004: 'Meta-ethics and the Problem of Creeping Minimalism'. In Hawthorne 2004, pp. 23–44.
- Egan, Andy, John Hawthorne, and Brian Weatherson 2005: 'Epistemic Modals in Context'. In Preyer and Peter 2005, pp. 131–68.
- Egan, Andy and Brian Weatherson (eds) 2011: *Epistemic Modality*. Oxford: Oxford University Press.
- von Fintel, Kai and Anthony Gillies 2007: 'An Opinionated Guide to Epistemic Modality'. In Gendler and Hawthorne 2007, pp. 32–62.

⁶¹ Thanks to David Charles, Sarah Moss, Timothy Williamson, Seth Yalcin, the participants in the 2011 Oxford/Paris semantics workshop, the editor of *Mind*, and four anonymous referees. Special thanks to Emanuel Chemla and Daniel Rothschild, whose empirically informed reply to an earlier version of this paper prompted substantial improvements.

- Fox, Danny 2007: 'Free Choice and the Theory of Scalar Implicatures'. In Sauerland and Stateva 2007, pp. 537–86.
- Gabbay, Dov and Franz Guenther 2002: *Handbook of Philosophical Logic*, vol. 8, second edition. Dordrecht: Kluwer.
- Gendler, Tamar Szabó and John Hawthorne (eds) 2007: *Oxford Studies in Epistemology*, vol. 2. Oxford: Oxford University Press.
- Gillies, Anthony 2004: 'Epistemic Conditionals and Conditional Epistemics'. *Noûs*, 38, pp. 585–616.
- 2010: 'Iffiness'. *Semantics and Pragmatics*, 3, pp. 1–42.
- 2011: "'Might" Made Right'. In Egan and Weatherson 2011, pp. 108–30.
- Glanzberg, Michael 2005: 'Focus: A Case Study on the Semantics–Pragmatics Boundary'. In Szabó 2005, pp. 72–110.
- Grice, H. P. 1989: *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hacking, Ian 1967: 'Possibility'. *Philosophical Review*, 76, pp. 143–68.
- Hawthorne, John (ed.) 2004: *Philosophical Perspectives 18: Ethics*. Oxford: Blackwell.
- 2007: 'Eavesdroppers and Epistemic Modals'. *Philosophical Issues*, 17, pp. 92–101.
- Iatridou, Sabine 2000: 'The Grammatical Ingredients of Counterfactuality'. *Linguistic Inquiry*, 31, pp. 231–70.
- Karttunen, Lauri and Stanley Peters 1979: 'Conventional Implicature'. *Syntax and Semantics*, 11, pp. 1–56.
- Klinedinst, Nathan and Daniel Rothschild 2012: 'Connectives Without Truth-Tables'. *Natural Language Semantics*, 20, pp. 137–75.
- Kolodny, Niko and John MacFarlane 2010: 'Ifs and Oughts'. *Journal of Philosophy*, 107, pp. 115–43.
- Kratzer, Angelika 2012: *Modals and Conditionals*. Oxford: Oxford University Press.
- MacFarlane, John 2011: 'Epistemic Modals are Assessment-Sensitive'. In Egan and Weatherson 2011, pp. 144–78.
- Preyer, Gerhard and Georg Peter 2005: *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Oxford University Press.
- Recanati, François 2002: 'Unarticulated Constituents'. *Linguistics and Philosophy*, 25, pp. 299–345.
- Roberts, Craige 1989: 'Modal Subordination and Pronominal Anaphora in Discourse'. *Linguistics and Philosophy*, 12, pp. 683–721.

- Rooth, Mats 1992: 'A Theory of Focus Interpretation'. *Natural Language Semantics*, 1, pp. 75–116.
- Sauerland, Uli 2004: 'Scalar Implicatures in Complex Sentences'. *Linguistics and Philosophy*, 27, pp. 367–91.
- Sauerland, Uli and Penka Stateva (eds) 2007: *Presupposition and Implicature in Compositional Semantics*. New York: Palgrave Macmillan.
- Stanley, Jason 2000: 'Context and Logical Form'. *Linguistics and Philosophy*, 23, pp. 391–434.
- Sorensen, Roy 2009: 'Meta-agnosticism: Higher-Order Epistemic Possibility'. *Mind*, 118, pp. 777–84.
- Stalnaker, Robert 1975: 'Indicative Conditionals'. *Philosophia*, 5, pp. 269–86.
- Stephenson, Tamina 2007: 'Judge Dependence, Epistemic Modals, and Predicates of Personal Taste.' *Linguistics and Philosophy*, 30, pp. 487–525.
- Spector, Benjamin 2007: 'Scalar Implicatures: Exhaustivity and Gricean Reasoning'. In Aloni, Butler, and Dekker 2007, pp. 225–50.
- Szabó, Zoltan Gendler (ed) 2005: *Semantics versus Pragmatics*. Oxford: Oxford University Press.
- Williamson, Timothy 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.
- Yalcin, Seth 2007: 'Epistemic Modals'. *Mind*, 116, pp. 983–1026.
- 2011: 'Nonfactualism about Epistemic Modality'. In Egan and Weatherson 2011, pp. 295–332.
- 2012a: 'Context Probabilism'. In Aloni et al. 2011, pp. 12–21.
- 2012b: 'A Counterexample to Modus Tollens'. *Journal of Philosophical Logic*, 41, pp. 1001–24.