

10

Nonfactualism about Epistemic Modality

Seth Yalcin

1. Introduction

When I tell you that it's raining, I describe a way the world is—viz., rainy. I say something factual, something whose truth turns on how things are with some aspect of the world. Likewise when I tell you that *the weatherman thinks* that it's raining. Here the truth of what I say turns on a different feature of the world, namely, the weatherman's state of mind. Likewise when I tell you that *I think* that it's raining. Here the truth of what I say turns on yet another feature of the world, namely, *my* state of mind.

Nothing like tedious platitudes to set the mood.¹ Okay—what about when I tell you that it *might* be raining? Or that it is *probably* raining? Or that it *must* be raining? In these cases, am I again to be understood as describing a way the world is?

An affirmative answer would be nice. For it seems like it would mean less work. It would let us take the view that sentences like these—sentences with epistemic modal operators taking wide scope—are not special. It would let us apply to these sentences whatever semantic and pragmatic explanatory strategies we already apply to other uncontroversially descriptive, fact-describing discourse. And, from a distance at least, an affirmative answer seems anyway not hard to pull off. Epistemic modals are so-called, after all, because they seem to serve to communicate information about some epistemic state or state of evidence. And states of evidence are, of course, aspects of the world. One could try, then, understanding epistemically modalized sentences—these sentences about what might or must be, or about what is probable—as telling

I am indebted to audiences at MIT, Berkeley, Princeton, NYU, and Arché for helpful feedback. Thanks especially to Andy Egan, John MacFarlane, Dilip Ninan, Robert Stalnaker, and Stephen Yablo for illuminating conversations on these matters. An earlier draft of this paper appeared in Yalcin (2008).

¹ Of course, one philosopher's tedious platitude is another's controversial thesis. An eliminativist about mental content would be uncomfortable with my opening, as would be a philosopher who takes content ascriptions to be normative, and who takes normative discourse not to be straightforwardly factual. I must set these views aside here.

how things are with some epistemic state or other, or with some body of evidence or other, in the world.

Since an affirmative answer means less work and looks not hard to pull off, little wonder that that answer is a popular one. Indeed, it has some title to being called the *standard view* about epistemic modality in philosophy. (Or at least, it had this title until recently.) One might spin the standard view either as a metaphysical thesis or as a linguistic thesis. We could call the linguistic thesis *descriptivism* about the content of epistemic modal discourse. To a rough first approximation, descriptivism is the idea that epistemic modal talk serves fundamentally to describe some feature of reality, to say how some aspect of the world is. (Compare the uncontroversial thesis of descriptivism about the content of weather talk.) The metaphysical thesis is *factualism* about epistemic modality. To a rough first approximation, factualism about epistemic modality is the view that there is a certain class of facts, the facts about what is (epistemically) possible, or probable, or necessary. For these facts to obtain is a matter of the world being one way rather than another. (Compare the uncontroversial thesis of factualism about the weather.) Descriptivism about the content of epistemic modal talk and factualism about epistemic modality are, I take it, really the same thesis in slightly different keys. I will use both terms for the view.

In this paper I am interested in developing an alternative to the factualist picture. With caveats to be provided in due course, the positive account I will set out could plausibly be called a kind of *expressivism* about epistemic modal discourse, and a kind of *nonfactualism* about epistemic modality. Before I begin developing this alternative, however, more should be said about the sort of view it is an alternative to. In the next section I say more about how I want to understand the descriptivist, factualist picture of epistemic modality. I then turn to some problems for the view, problems which recommend investigation into nondescriptivist alternatives.

2. Factualism

Descriptivism, I already said, is the view that epistemic modal talk serves to describe reality. Let me clarify ‘epistemic modal talk’ and ‘serves to describe reality’.

By ‘epistemic modal talk’, I have in mind foremost sentences that are modalized with natural language epistemic modal operators. For instance, *It is possible that it is raining*, *It might be raining*, *It could be raining*, *It is probably raining*, *It is likely that it is raining*, and *It must be raining* all have readings on which the modals they contain are interpreted epistemically. (With *might*, *likely*, and *probably*, the epistemic reading is the preferred reading, if not the only reading; with *could*, *possible*, and *must*, other readings, such as a deontic reading, are often possible.) I don’t attempt an operational definition of this class of modals now; the ultimate project is to provide a theory which delimits the class more precisely. Only let me be explicit that by ‘epistemic modal operator’, I don’t have in mind complex operators such as ‘for all I know, it might be

that’—operators with simple epistemic modals scoped under epistemic attitude verbs. The importance of excluding these complex operators will become clear later. I will also avoid interactions with tense, restricting myself to the case where these modals take apparently present-tensed complements.

By ‘serves to describe reality’, I mean that epistemically modalized sentences serve to represent the world, or one’s situation in the world, as being a certain way. Relative to context, the content of the sentence determines, and is understood as determining, a condition on metaphysically possible worlds or situations. It has the effect of dividing the space of possible ways things might be into those which conform, and those which fail to conform, with how things are represented as being; and moreover the fact that the sentence effects this division forms a crucial part of the explanation for its communicative import. Let me call a rule for dividing the space of possible worlds or situations *factualist truth conditions*. A descriptivist provides factualist truth conditions for epistemic modal talk.

This is close to what we want, but my characterization of descriptivism has so far blurred over an important distinction. This is the distinction between the *compositional semantic contribution* of a declarative sentence, and what we might call the *informational content* of the sentence when it is tokened in a context. Though often conflated in practice, these notions correspond to distinct theoretical roles. The former notion is for use in articulating our tacit semantic competence with the language—in particular, our in-principle ability to interpret, via a finitely specifiable competence, an infinitude of sentences. In contrast, the notion of informational content—what I have in mind here also sometimes gets called *the proposition expressed* or *what is said*—is best construed as a notion at the semantics-pragmatic interface. It comes into play when we ask:

What information does a given sentence normally communicate in virtue of its compositional semantics, relevant features of the context of utterance, and the standing pragmatic norms which are common knowledge among speakers of the language?

To answer this kind of question, one needs some additional theory. Most obviously, one needs a general conception of what information or content is. One needs a theory of content. Such a theory might be influenced, or even largely driven by, considerations external to formal semantics *per se*.² One needs also some pragmatic theory—in particular, some view about what conventional rules govern the linguistic communication of information. The notion of informational content is constrained from these directions. While we can expect that knowledge of the semantic value of a sentence should be a fundamental part of the explanation of how a listener recovers the informational content of the sentence in context, it is not necessary to assume that the semantic value of a sentence is literally *identical* to the informational content of the sentence in context. As Lewis (1980) has noted, what we need is only the idea that the informational content of a sentence in context can be recovered in some

² Theories of content wherein the formal semantics of natural language plays no essential role are quite familiar in philosophy; Dretske (1981) and Stalnaker (1984) are two well-known examples.

systematic way from its semantic value (together with context and whatever standing, commonly known pragmatic principles there may be).

With the conceptual distinction between these two notions in mind, we can make it official that descriptivism, as I want to understand it here, is a view about the informational content of an unembedded, declarative epistemic modal sentence. It is the view that the informational content of these sentences has the effect of dividing the space of possible ways things might be into those which conform, and those which fail to conform, with how things are represented as being; and moreover it is the view that the fact these sentences effect such a division forms a crucial part of the explanation for their communicative import. And the thought is that in these respects, epistemic modal talk is just like ordinary, prosaically factual talk. So understood, the formal semantics of epistemic modals of course bears on the thesis of descriptivism, but its bearing is indirect. The question of whether descriptivism is true is not equivalent to the question of the compositional semantics of epistemic modal operators.

I hope descriptivism sounds like the straightforward view it is. Some examples of descriptivist views will help to round out the picture. Start with descriptivism about epistemic possibility talk. G. E. Moore writes:

People *in philosophy* say: The propositions that I'm not sitting down now, that I'm not male, that I'm dead, that I died before the murder of Julius Caesar, that I shall die before 12 tonight are 'logically possible'. But it's not English to say, with this meaning: It's possible that I'm not sitting down now etc.—*this* only means 'It is not certain that I am' or 'I don't know that I am'. (1962: 184)

Moore's view is descriptivist simply because according to it, epistemic possibility sentences in context are descriptions of the epistemic state of some agent in the world.

Most descriptivists agree with Moore's basic idea, that these sentences serve to describe the epistemic situation of some agent or agents. The internal debate among descriptivists concerns the detailed nature of the description—for instance, which agents matter, or what aspects of the agents' evidential situation are relevant. For example, three or four epicycles of analysis down from Moore, DeRose proposes that

S's assertion "It is possible that P" is true if and only if (1) no member of the relevant community knows that P is false, and (2) there is no relevant way by which members of the relevant community can come to know that P is false (1991: 593–4)

while Stanley seems to suggest that these sentences describe the epistemic state of some contextually given knower *A*:

It is possible_{*A*} that *p* is true if and only if what *A* knows does not, in a manner that is obvious to *A*, entail not-*p*. (2005: 128)

The particular motivations for these departures from Moore's position needn't detain us. The point is just that, although Moore, DeRose, and Stanley all differ on exactly what facts epistemic possibility sentences describe, they all agree that these sentences serve to describe some facts or other, some feature of the world.

Those are examples of descriptivism about epistemic possibility. What about epistemic necessity? If, as is widely assumed, epistemic necessity modals (\Box) are the logical duals of epistemic possibility modals (\Diamond) in the sense that

$$\Diamond\phi \leftrightarrow \neg\Box\neg\phi$$

then we can expect that each of the above accounts of epistemic possibility straightforwardly generates an account of epistemic necessity. So, given duality, Moore's view would be that 'It must raining', on the epistemic reading, is true just when 'I know it's raining' is; and so on for the other two views. It should be clear that the resulting positions on epistemic necessity are no less descriptivist than the positions on epistemic possibility they are constructed from. The duality of epistemic possibility and necessity is plausible; I will assume it throughout.

(But let me warn in advance: epistemic necessity will play second fiddle to epistemic possibility in this paper. What I say about it will be driven mostly by considerations about epistemic possibility, and the assumption of duality.)

Last, probability operators such as 'probably' and 'it is likely that', which I will abbreviate as ' Δ '. Here, a simplistic example of a descriptivist position can be abstracted from the Bayesian paradigm: for one to say 'It's probably raining' is for one to say that one's credence in rain is above one-half, or above some contextually-determined value.³ In calling something 'probable', one describes one's credal state. A second position, closer in form to the descriptivist accounts of (non-probabilistic) epistemic modality just described, adverts to some tacit body of knowledge or evidence. Suppose a body of evidence induces, or is representable by, a probability measure over a domain of propositions. Then instances of $\Delta\phi$ can be understood to say that the proposition that ϕ has some highish value according to the measure induced by the body of evidence determined by the context in which the sentence is uttered.⁴ They would, in short, be factual claims about some contextually determined body of evidence.

The factualist, descriptivist view requires no fundamentally new assumptions. Epistemic modal sentences are taken to determine possible worlds truth conditions, so from a semantic point of view we can compositionally supply epistemic modal clauses with truth conditions in ordinary fashion. (The leading semantics is Kratzer's: see Kratzer (1977, 1981, 1991b); see also Lewis (1979).) Pragmatically, too, we needn't make waves. We can retain a familiar picture of communication, a picture that gives no special place to epistemic modal talk. Whether I say that it is raining or I say that it is probably raining, the story about what is happening, at least in straightforward cases, can be the usual Gricean one: I believe myself to have some information about the world, and wish to impart it to you; I say something whose truth turns on whether

³ Jeffrey seems to have something like this in mind when he writes: "If you say the probability of rain is 70% you are reporting that, all things considered, you would bet on rain at odds of 7:3" (2004: 3). (For Jeffrey, one's credence just is a matter of one's disposition to bet.)

⁴ To ease exposition, I will be loose about use and mention in this paper.

this information is true, presuming common knowledge of the language; in so doing I intend for you to come to accept that information, acting with the expectation that my intention to communicate that information is mutually recognized. More needs to be said to fill in the details, of course; but suffice to say that, for the descriptivist, the details will be filled in just the same ways for epistemic and non-epistemic modal talk alike.

Because descriptivism makes no semantic or pragmatic waves, there is a certain presumption in favor of it. In the next section I attempt to undermine this presumption.

3. Challenges for Factualism

3.1. *Epistemic contradictions*

Notice that the following sentences sound awful.⁵

- (1) # It's raining and it might not be raining.
- (2) # It's raining and it probably isn't raining.

Let us call sentences like these—sentences of the schematic form $(\phi \wedge E\neg\phi)$, where E is an epistemic modal—*epistemic contradictions*. Why do epistemic contradictions (1) and (2) sound awful?

At first glance, a descriptivist-friendly explanation seems easy enough. A descriptivist might try saying that these sentences sound terrible because, thanks to the semantics of the epistemic modals, these sentences both truth-conditionally entail

- (3) # It's raining and I don't know its raining.

relative to context. Obviously, (3) is Moore-paradoxical. Therefore, says the descriptivist, (1) and (2) should be pragmatically defective in whatever way (3) is defective. The defect in (1) and (2) is parasitic, as it were, on (3). Epistemic contradictions are not contradictions in any semantic sense. They are just Moore-paradoxical sentences in new guise.

The situation is not so simple, however. (1) and (2) are more than merely pragmatically defective. The conjuncts in these sentences are incompatible in a more robust sense. We can see this when we attempt to embed these conjunctions into larger constructions. Consider, for instance, the imperatives:

- (4) # Suppose it's raining and it might not be raining.
- (5) # Suppose it's raining and it probably isn't raining.

These imperatives sound incoherent and self-defeating. The fact that they do not make sense is not explained by the assumption that the conjunctions they each embed both truth-conditionally entail (3), because (3) is perfectly easy to embed under 'suppose':

⁵ Here I review some points made in more detail in Yalcin (2007).

(6) Suppose it's raining and I don't know its raining.

(Indeed, the *intelligibility* of sentences like (6) is a popular motivation for denying that classically Moore-paradoxical sentences are contradictions in any semantically rich sense.) Or again, epistemic contradictions never sound acceptable in the antecedent position of an indicative conditional:

(7) # If it's raining and it might not be raining, then . . .

(8) # If it's raining and it probably isn't raining, then . . .

Conditionals that begin in this way seem beyond repair. But Moore-paradoxical sentences are acceptable in this environment:

(9) If it's raining and I don't know it, then I will get wet.

Compare that with the nonsensical

(10) # If it's raining and it might not be raining, then I will get wet.

The conditional (10) is particularly telling. If it really were the case that, relative to context, 'It might not be raining' entailed 'I don't know that it's raining', we would expect (10) to be about as acceptable as (9). But the difference in acceptability could hardly be greater.

The upshot is this. Epistemic contradictions 'project their unacceptability', as it were, in the embedded contexts described above. Moore-paradoxical sentences do not. The defect in sentences which embed epistemic contradictions is therefore not parasitic on Moore's paradox. It must be explained in some other way.

And the problem is that it is not clear how to explain it plausibly along descriptivist lines. Descriptivists want to tell us that epistemic contradictions such as (1) and (2) above have factulist truth conditions. In particular, they want to tell us that these sentences have *non-empty* factulist truth conditions, truth conditions that obtain in some possible situation.⁶ If the truth conditions of (e.g.) 'It isn't raining and it might be raining' are non-empty, however, it seems there should be nothing at all preventing us from hypothetically entertaining the obtaining of these conditions. But we can't, evidently; there is no coherent way to entertain the thought that it isn't raining and it might be raining. This is quite unexpected from a descriptivist perspective.

Let me be clear about the extent to which these embedding facts do and do not tell against the descriptivist view. Considerations about where exactly epistemic modal clauses can felicitously embed most directly constrain the compositional semantics of epistemic modals. As we have noted, however, compositional semantics is one thing, and informational content another. The descriptivist therefore has a recourse: she could attempt to exploit the gap between these two notions in covering the above facts. She could do this if she can find semantic values for epistemic modals (and for attitude verbs like 'suppose', and for indicative conditionals) that will lead us to

⁶ I assume the descriptivist does not wish to maintain that the conjuncts of an epistemic contradiction are truth-conditionally incompatible, hence that 'It might be raining' truth-conditionally entails 'It's raining'.

expect defectiveness in the case of the problematic embedded epistemic contradictions above, but which nevertheless contribute to determining a coherent (albeit Moore-paradoxical) informational content for the unembedded epistemic modal sentences (1) and (2).

As far as I can see, the descriptivist must take this sort of recourse in order to make sense of the embedding facts. That makes for a somewhat complicated wrinkle in the descriptivist view. But since a wrinkled view along these lines is certainly logically possible,⁷ the embedding facts just described do not straightaway refute descriptivism. They do suggest, however, that descriptivism has the following surprising feature: when one says that it might be raining, the informational content one expresses is not the content one is related to when one supposes it might be raining. Moreover if, as seems plausible, the semantics of ‘suppose’ is taken to relevantly pattern with other attitude verbs such as ‘believes’, so that we can explain what is troubling about

(11) John believes it’s raining and it might not be raining.

along structurally similar lines, the descriptivist view can be expected to have the result that to believe that it might be raining is not to believe the proposition (informational content) one would express when one says that it might be raining (that is, utters the unembedded declarative sentence ‘It might be raining’). This result is surprising. Ultimately, I will offer a view that avoids this result.

3.2. *Assertability and disagreement*

We sometimes disagree, not merely about what is the case, but also about what might be the case, and about what is probably the case. The second concern about descriptivism is that it is hard to see how to deliver factalist truth conditions for epistemic modal talk which make sense of this kind of disagreement.

The trouble appears to have been first noted by Huw Price. Price considers the idea of assigning $\Delta\phi$ factalist truth conditions along the lines of ‘Given the existing evidence, it is probable that ϕ ’. He observes that the phrase ‘the existing evidence’ is ambiguous, admitting a spectrum of readings from the more subjective to the more objective. He first attempts a subjective reading of the phrase, along the lines of ‘the evidence of which I [the speaker] am actually aware’. He objects that:

If I disagree with your claim that it is probably going to snow, I am not disagreeing that given *your* evidence it is likely that this is so

and contrariwise:

Indeed, I might *agree* [with your claim] that it is probably going to snow and yet think it false that this follows from your evidence. (Price 1983: 404)

Here the problem is that the proposed truth conditions for ‘It is probably going to snow’ are too weak to make sense of cases both of appropriate agreement and appropriate disagreement.

⁷ The ‘diagonal view’ discussed in Yalcin (2007) is a precise version of such a view.

Next he tries an objective reading of ‘the existing evidence’, along the lines of ‘the evidence accessible in principle’. He objects that this more objective reading wouldn’t square with the facts about when $\Delta\phi$ is felicitous to assert:

consider the surgeon who says, ‘Your operation has probably been successful. We could find out for sure, but since the tests are painful and expensive, it is best to avoid them.’ The accessibility, in principle, of evidence which would override that on which the [probability] judgment is based, is here explicitly acknowledged. (Price 1983: 405)

Here the surgeon says $\Delta\phi$, but leaves open whether ϕ is probable given the evidence accessible in principle. No surprise he would leave *that* question open, after all; he simply does not have the evidence accessible in principle. Hence his statement of $\Delta\phi$ is not well understood as speaking to a question about the evidence accessible in principle.

This now provokes the question: what or whose evidence is relevant to settling the truth of a given claim of $\Delta\phi$? We appear to need something in between the evidence of the speaker and the evidence available in principle. It is hard to see, however, how something in between could ever really be assertable for the speaker. Something in between, after all, is by definition beyond the scope of the speaker’s evidence. If we settled on something in between, our speaker would still be pictured as saying something whose truth turns on a body of evidence that she does not have. It seems she would be pictured as speaking, and knowingly speaking, from a position of ignorance, making a stronger claim than is warranted by her evidence alone. Her speech act looks in danger of being irrational.⁸

The objection to descriptivism, then, is that it faces a tension. Either descriptivist truth conditions systematically fail to capture the truth-value judgments that people actually make (by being too weak to capture the disagreement facts), or it captures these judgments but turns users of epistemic modal sentences into irrational asserters (by picturing them as making claims about/from a body of evidence they don’t have). The difficulty here recurs exactly with epistemic possibility claims, as the reader may confirm by replacing ‘probably’ with ‘possibly’ in Price’s examples.

3.3. *Conflicting intuitions*

Closely related, a third problem with descriptivism is that it leads us to expect clear intuitions in cases where intuitions are not clear. Consider the following case:⁹

⁸ Couldn’t her evidence include information about the state of some other, not-yet-possessed body of evidence? And in that case, couldn’t her evidence provide her with the warrant to make claims about what is made probable by this not-yet-possessed evidence? But it is unclear whether one can have evidence that some other, larger body of evidence makes p probable without one’s own evidence itself making p probable.

⁹ These are usually called *eavesdropping cases*. See Egan et al. (2005), Egan (2007), von Fintel and Gillies (2008), and MacFarlane (this volume) for discussion. Some of these authors take the speaker judgments about eavesdropping cases to be less ambivalent than I do—on this matter we take different positions on what is ultimately an empirical question—and they use these cases to motivate different versions of *relativism* about epistemic modal talk. I will present a different view about how to respond to these cases.

Fat Tony secretly plants highly compelling evidence of his murder at the docks. The evidence is discovered by the authorities, and word gets out about his apparent death. The next evening, from his safehouse, Fat Tony watches a panel of experts on the news discussing the situation.

Expert *A* has had a good look at the evidence found at the scene. “Fat Tony is dead,” he says.

Expert *B* has also had a good look at the evidence, but his assessment is more cautious. “Fat Tony might be dead,” *B* says.

We can all agree that Expert *A*, however reasonable his speech act was in light of the information available to him, spoke falsely. Things are not as he says they are. Okay; what about Expert *B*? Is what he said true or false? Let me remind you that Fat Tony’s planted evidence was highly compelling. Let me remind you also that Fat Tony is *definitely not dead*. And, before you settle on an answer, let me ask you also to ponder whether Fat Tony himself should agree with your answer.

Now, what do you say about what *B* said—true or false?

It appears that, as a matter of empirical fact, intuitions are unclear about cases such as this—cases where an epistemic modal claim is assessed for truth from outside the discourse context.¹⁰ Some are inclined to say that *B* spoke truly; others are inclined to say that *B* spoke falsely; everyone else shrugs, or proposes to change the question to one with a clearer answer.

What needs explaining for eavesdropping cases such as this, I think, is not any univocal intuition we all have about the epistemic modal claim made in the case. In the absence of a systematic empirical study, there seems to be no single intuition there to explain. Rather, what needs explaining is the absence of agreement, by competent speakers of English, on what the right answer is. What needs explaining are the conflicting intuitions. My point for now is just that conflicting intuitions are not expected on descriptivist assumptions. If *B*’s utterance is in the business of representing the world as being a certain way, as *A*’s presumably is, then either the world is that way, or it isn’t. Other things being equal, we’d expect intuition concerning the truth of *B*’s utterance to be about as clear as it is with *A*’s. Descriptivists have work to do, then, explaining why things are not equal.

A descriptivist might reply that this work is not really so hard. “Epistemic modals are, after all, highly context-sensitive on our view. Perhaps the lack of uniformity in judgments here is simply due to the fact that subjects considering this case resolve this context-sensitivity in different ways.”

But this reply is unsatisfactory. If the interpretation of epistemic modals is as context-sensitive as suggested, presumably there is at least one reading of the modal according to which what *B* says above is both (a) true and (b) assertable for *B*. (Perhaps a reading along the lines of ‘The evidence in the reach of *B* leaves open the possibility that Fat Tony is dead.’) Where multiple interpretations of a speaker’s utterance are

¹⁰ I say this not from the armchair, but from having surveyed 128 subjects on the matter. For a brief discussion, and details on the data see Yalcin and Knobe (2010).

possible, subjects tend to gravitate towards true and assertable readings, for the simple reason that true and assertable readings tend to make most sense of what the speaker is doing—they tend to be easier to situate into a rational overall pattern of action. But this would lead us to expect a fairly robust judgment that what *B* says is true, the incorrect result.

It is striking, incidentally, that the body of evidence allegedly relevant to assessing the truth of an epistemic modal claim should be so obscure to speakers who actually use these sentences. If these sentences really do advert to some tacit body of evidence, as standard versions of factualism maintain, why are we competent speakers of the language not able to articulate what this body is? This opacity is puzzling. It is not a feature of context-sensitive language in general. For instance, when we use quantifiers in ordinary discourse, typically a restriction on the quantifier is provided tacitly by context. But with sufficient description of context, speakers can typically recover what the intended restriction is; and where context is insufficient, speakers can typically indicate what further information is needed to settle the question. In contrast, the interpretation of epistemic modals seems, from a contextualist perspective, to be far less constrained. Appeals to the context-sensitivity of epistemic modals seem to be of questionable explanatory power here, then.

Let me summarize. We have accumulated three desiderata for a theory of the content of epistemic modal claims. Such a theory should:

- i. Be compatible with a plausible explanation of why **epistemic contradictions** are unembeddable.
- ii. Make sense of the **assertability and disagreement** facts concerning epistemic modal claims in context.
- iii. Make sense of the **conflicting intuitions** concerning epistemic modal claims in eavesdropping cases.

The first desideratum is plausibly understood as a constraint on the formal semantics of epistemic modals (together with the semantics of the relevant embedding environments). The second two desiderata are plausibly understood as constraining the pragmatics of epistemic modal claims—more precisely, their communicative content. The first of these pragmatic desiderata concerns the intra-contextual facts about how we assess epistemic modal claims *qua* participants in the discourse. The second concerns the extra-contextual facts about how we assess epistemic modal claims *qua* onlookers from outside the discourse.

Descriptivism is not well-positioned to satisfy these desiderata. It is time to take steps towards an alternative.

4. States of Mind

If we want to understand what is going on with epistemic modal talk, we may be better served by taking a less direct approach. Let us take a step back from the linguistic

facts and from direct questions about the truth conditions of epistemic modal clauses. Let us ask instead:

What is it to be in a state of mind which accepts what an epistemic modal claim says?

I will suggest that descriptivism rests on a mistaken answer to this question, and that getting the answer right is the first step towards clarifying the meaning and role of epistemic modal discourse. The focus of this section will be on developing a model for what it is to believe that something might be so, or that something is possibly so.

In stepping back from direct questions about the truth conditions of epistemic modal clauses and asking instead about what it is to be in an epistemically modal state of mind, I make a move characteristic of meta-ethical expressivists like Gibbard (1990, 2003). Rather than asking after the truth conditions of normative sentences directly, Gibbard asks first what it is to be normatively opinionated—that is, what it is to be in a state of mind with normative content. Once he develops a theory of this state of mind, he moves from that theory back to questions about the semantics of normative discourse and about the metaphysics of normativity. When direct inquiry into truth conditions does not seem to bear fruit, this kind of strategy—we could call it the method of *psychological ascent*, in contrast with the truth-conditions-focused method of semantic ascent—is, I think, sensible. I take this kind of strategy here.

We begin with epistemic possibility. I believe that it is possible that Bob is in his office; Frank believes that it might be raining in Topeka. What kind of states of mind are we each in? Doxastic states of mind, trivially. How to model a doxastic state of mind? For our purposes, we may represent a doxastic state by its informational content, abstracting for now from its functional role in cognition and action. How, then, to represent the informational content of a doxastic state of mind?

Start with a familiar picture of informational content in general. Information is foremost that which eliminates possibilities. To gain information is to transition to a state of mind which leaves fewer possibilities open as candidates for actuality. As a first approximation, then, let us represent a body of informational content as a set of possibilities, those possibilities left open by that informational content. So a state of belief is representable by a set of possibilities: intuitively, those not excluded by what is believed. The propositions true at each world in the set are the propositions believed by the agent. (Propositions, too, we will model via their truth-conditional content: a proposition is a set of possibilities, intuitively the possibilities with respect to which the proposition is true.)

As everyone knows, this classic possible worlds representation of belief faces acute problems.¹¹ Let me note now then that dialectically this classic picture will be serving as my point of departure, not arrival. Soon we will work this classic picture into

¹¹ For example, Frege's puzzle and the problem of logical omniscience (to name the two most commonly cited difficulties). The latter problem is discussed further below.




Fig. 10.1. $B_A \diamond \phi$: The descriptivist model.

something more realistic. Meanwhile this model, idealized as it is, will supply us with a useful starting point.

Equipped with this representation, we can provide an abstract picture of the descriptivist model of epistemic possibility beliefs—of what, according to the descriptivist, it is to believe that it is possible that Bob is in his office, or that it might be raining in Topeka.¹² The picture is very simple (Fig. 10.1).

The rectangle is logical space, the space of maximally specific metaphysical possibilities. A subset of those possibilities is the proposition that $\diamond\phi$, here the set of possibilities contained within the dashed ellipse. A believes that $\diamond\phi$ just when A 's belief worlds are a subset of the proposition that $\diamond\phi$. Thus for me to believe that Bob might be in his office is for a certain proposition—whatever proposition it is the descriptivist says it is—to be true throughout my belief worlds. Again, standardly the descriptivist's truth conditions are propositions about some body of evidence, where this body of evidence includes the knowledge of the agent doing the believing. As a result, the typical descriptivist picture is one according to which states of $\diamond\phi$ -belief are *second-order* states of mind, states of belief about (perhaps *inter alia*) one's state of knowledge. And this prompts the question: when I believe Bob might be in his office, am I in a second-order state of mind?

We could try asking it like this. Is the question, "Why believe Bob might be in his office?" in part the question, "Why believe that I don't know that Bob isn't in his office?" Pre-theoretically, the idea seems to have little to recommend it. Our initial question seems to be about Bob's location, not about my views about Bob's location. The question "Why believe Bob might be in his office?" seems instead equivalent to the question, "Why fail to believe that Bob isn't in his office?" This latter question

¹² I have in mind here the 'unwrinkled' descriptivist, the descriptivist who holds that in saying that it might be raining, one expresses the proposition one believes when one believes that it might be raining.

is clearly not a question about what to believe about one's knowledge. It is just a question concerning what to believe about where Bob is.¹³

Other considerations lead in the same direction. Suppose we are eating dinner, and my dog Fido comes into the room and heels by my chair. Occasionally I toss Fido a bone at dinner, but usually I don't. You ask why Fido is sitting there staring at me. I say:

(12) Fido thinks I might give him a bone.

What I say speaks to your question. Unless you have a particular theory of epistemic modals, I doubt you would flinch at this remark. But what exactly am I saying? Does my remark in part mean, as standard versions of descriptivism would recommend, that Fido believes that it is left open by what he knows that I will give him a bone? That is a bit much. The truth of (12) does not turn on *recherché* facts about canine self-awareness. Surely (12) may be true even if Fido is incapable of such second-order states of mind.

These considerations suggest that the question of whether $\diamond\phi$ is 'transparent', as it were, to the question of whether ϕ . I think this is reflected in the kinds of reasons we understand to support epistemic possibility beliefs. Naively, correctly believing that $\diamond\phi$ is a matter of there being an absence of conclusive reason to believe that $\neg\phi$. Correctly believing that ϕ is a matter of there being conclusive reason to believe that ϕ . Both kinds of reason concern how to settle one's doxastic state toward the proposition that ϕ . Believing that ϕ and believing that $\diamond\phi$ are states of mind supported by reasons of the same category.

It seems, then, that we have found another desideratum for a theory of the content of epistemic modal claims. Such a theory should:

- iv. Avoid the assumption that to believe that something might be the case is to be in a **second-order state of mind** (i.e. a state of belief *inter alia* about one's own state of mind).

This is another desideratum that descriptivism is not well placed to capture. If one expresses a proposition one believes when one says (e.g.) 'It might be raining', and that proposition has the epistemic-state-describing truth conditions assigned to it by standard versions of descriptivism, it is a very short step to the thought that to believe it might be raining is to believe that very proposition.

Ask now: what minimal modification to the descriptivist model would be required to satisfy this new desideratum (iv)? I suggest that the modification is this one depicted in Fig. 10.2.

¹³ While, I think, suggestive, these observations are perhaps not decisive. That is because it can be difficult to disentangle, from a first-person point of view, questions about what the world is like from questions about what one believes the world is like. As Evans observed, "If someone asks me 'Do you believe that there will be a third world war?', I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question 'Will there be a third world war?'" (1983: 225). (Though the point should not be overstated. The questions, "Why believe that ϕ ?" and "Why believe that you believe that ϕ ?" need not always have the same answer.)




Fig. 10.2. $B_A \diamond \phi$: The first-order model.

On this revised model, due essentially to Frank Veltman,¹⁴ there is no proposition that $\diamond \phi$ at work. There are no ' $\diamond \phi$ -worlds'. The question of whether A believes that $\diamond \phi$ is just the question whether A 's belief worlds leave open possibilities wherein the proposition that ϕ is true. To believe Bob might be in his office is simply to be in a doxastic state which fails to rule out the possibility that Bob is in his office. It is a first-order state of mind. One might think of this as an 'adverbialist' model of epistemic possibility belief. Such beliefs do not correspond to a distinctive class of believed contents; rather, they correspond to a distinctive way of being doxastically related to a proposition. Note that on the first-order model, 'epistemic modal' is an unfortunate moniker, for there is no special role for a state of knowledge in this picture.

The first-order model is a considerable advance over the descriptivist model. It avoids the implausible idea that epistemic possibility beliefs are second-order states of mind, and in a way that lets us see why reasons that support belief that ϕ would be *ipso facto* reasons that support belief that $\diamond \phi$. Besides satisfying our newly uncovered fourth desideratum, the first-order model helps also with our first three desiderata.

I. Making sense of epistemic contradictions. There is no difficulty at all giving a formal semantics for epistemic modals, indicative conditionals, and attitude verbs which can cover the embedding facts about epistemic contradictions in a manner consistent with the first-order model. (Indeed, as noted in footnote 14, the first clear occurrence of the idea was within a formal semantics.) Veltman's dynamic semantics for epistemic modals, when supplemented with appropriate semantics for conjunction, attitude verbs (e.g. Heim 1992) and indicative conditionals (e.g. Gillies 2004)

¹⁴ See Veltman (1985, 1986, 1996), where this model is suggested by the dynamic semantics for epistemic modals developed in these works. (The first-order model is roughly what one would get as the truth conditions of epistemic possibility modal-embedding belief reports, were one to combine Veltman's semantics for epistemic modals with a straightforward dynamic possible worlds semantics for 'believes', such as Heim (1992).) Let me be clear that although what I am calling "the first-order model" is suggested by work in semantics, it is not itself a thesis in compositional semantics. Rather, it is a thesis about how to model a certain kind of state of mind.

will predict the defectiveness of epistemic contradictions in the relevant embedded contexts. (See Yalcin 2007 for more discussion, and for an example of a static semantics which can also accommodate the data.) Thus from a semantic point of view, the first-order model poses no special difficulties. On the contrary, it can be *motivated* by semantic considerations.

II. Making sense of the assertability and disagreement facts. We noted that the assertability and disagreement facts concerning unembedded epistemic modal claims are hard to explain under the assumption that there are some factual truth conditions which constitute their informational content. The first-order model suggests a different way of thinking about what we are up to when we say that something is possible, or might be the case. It recommends the idea that in modeling the communicative impact of an epistemic possibility claim, we construe the objective as one of coordination on a certain global property of one's state of mind—the property of being compatible with a certain proposition—not one of coordination concerning the way the world is.

First, this lets us avoid the demand, incumbent on the descriptivist, to say what or whose evidence is relevant to settling the truth of a given epistemic possibility claim. These claims do not have factualist truth conditions; *a fortiori* they do not have truth conditions turning on how things are with some body of evidence. Second, it lets us see the assertability and disagreement facts in a different light. To believe something might be the case, on the first-order view, is not really to embrace any positive thesis about how the world is. Rather, it is a way of lacking information; it is effectively a state of failing to believe something. And the relevant point to note here is that our intuitive notion of (dis)agreement is not really trained on such states of mind. If someone believes it's raining in Topeka right now, then you *agree* if you also think it's raining in Topeka right now, and you *disagree* if you think it's not raining in Topeka right now. If you neither agree nor disagree, you are agnostic on the matter. So understood, disagreement goes beyond mere failure of agreement. But what then is it to *agree with the agnostic* on this issue? The question feels ill-posed. This ill-posed question is, on the first-order view, similar to the question of what it is to agree or disagree with someone who believes it might be raining.¹⁵

If epistemic possibility claims do not have factualist truth conditions, what do they have? There are various ways one could build a formal model of the communicative impact of epistemic possibility claims—of their pragmatics—consistent with the first-order idea that the objective of such claims is to achieve coordination between states of mind on the openness of a possibility. One way is the one already standard in the

¹⁵ One might try saying that to agree with the agnostic is to just also be agnostic, but this line is not promising. For what then would it be to disagree with the agnostic? Is it to take a definite stand on the relevant proposition? But it is not plausible to say that, merely in virtue of my lack of a positive stand on myriad propositions, I thereby disagree with anyone and everyone who takes a stand on these propositions. Suppose I don't have a view about where you parked your car, and you do. I do not thereby disagree with you, in any interesting sense.

dynamic semantics literature on epistemic modals stemming from Veltman's work (see e.g. Beaver 2001). Building on Stalnaker (1970), we associate with each conversation a *context set*, a set of possible worlds left open by what is mutually presupposed by the discourse participants. The communicative impact of an utterance is then formally modeled in terms of its characteristic tendency to change or update the context set. With ordinary factual discourse, the tendency is to eliminate worlds from the context set—to rule out ways the world might be. But with epistemic possibility claims, the tendency is only to ensure coordination on the leaving of a certain possibility open. The context set is just a representation of the speakers' presuppositions, and to presuppose something might be the case is exactly like believing something might be the case: it is a matter of the relevant proposition being compatible with the content of that state of mind.

III. Making sense of conflicting intuitions about eavesdropping. We observed above that we tend to have conflicting intuitions about the truth value of an epistemic modal claim when we are outside the discourse context and in a better epistemic position (with the respect to the epistemically modalized proposition) than those within the discourse. When Expert *B* says 'Fat Tony might be dead' in the scenario envisaged earlier, it is not clear what truth value the claim deserves.

The reason it is not clear, I suggest, is that Expert *B*'s speech act does not serve to describe the world. There is no way the world could be, or could fail to be, which would settle the question of the truth of the sentence. For this sentence there is no answering the question,

- (T) Is the content of this speech act **true** in the sense that its factualist truth conditions characterize the actual world?

for it has a false presupposition. His utterance does not have factualist truth conditions.

The point of the speech act on the story I recommend is, again, to engender coordination among one's interlocutors with respect to the property of states of mind the sentence semantically expresses in context. Insofar as the claim has a content which is communicated, it is simply this property (and not a rule for eliminating possibilities—not a kind of factualist content). When assessing this kind of speech act for correctness, we cannot ask (T). At best it seems we ask one of two things:

- (R) Is the speech act **rational** in the sense that someone equipped with the evidence of the speaker would be responding appropriately to the evidence by accepting the content of the speech act?
- (A) Is the speech act **advisable** in the sense that a person equipped with full information about the relevant situation would be responding appropriately to that information by accepting the content of the speech act?¹⁶

¹⁶ Gibbard makes an analogous distinction concerning questions about what it 'makes sense' to do (1990: 18–19).

Given this distinction, it is easy to see that Expert *B*'s speech act was *rational* but *inadvisable*. His state of mind responded appropriately (in one sense) to the evidence, but it is not the state of mind we would recommend to him given our superior epistemic position.

When we are asked about the truth value of claim in a given context, typically we understand the question to be (T). But where the claim is epistemically modalized, that question cannot arise. We therefore look for other criteria to assess the sentence for correctness; and the two kinds of features we check for instead, I suggest, are rationality and advisability. When ordinary speakers are asked, 'Is what Expert *B* said true?' some of them interpret the question as (R), and they answer 'yes'. Others interpret the question as (A), and they answer 'no'. Still others feel the intuitive pull of both interpretations. These enlightened subjects reject the question and say: "Look: Expert *B* was right to say what he did, given what he knows. But if he were to say that to me, I'd reject it, because I know the facts of the case." These speakers tacitly recognize that, as far as the correctness of the speech act goes, we can ask either of (R) or (A), but that no further question (T) arises.

We have now effectively wedded the first-order view to a kind of nonfactualism about epistemic possibility and a kind of expressivism about the associated discourse. To believe something is possible is not to take the world to be one way rather than another; it is not to think a certain sort of fact obtains. It is for one's state of mind to have a certain global property, one not reducible to a condition on worlds. We approach the question of the communicative impact of epistemic possibility claims from this perspective. To say that a proposition is possible, or that it might be the case, is to express the compatibility of the proposition with one's state of mind,¹⁷ with the intention of engendering coordination on this property with one's interlocutor.

This expressivist, nonfactualist view about epistemic possibility is motivated by the facts—in particular, it is motivated by linguistic facts and by intuitive considerations about what it is to accept that something is possible. Unlike some versions of its meta-ethical cousin, this brand of nonfactualism is not driven by a metaphysical concern about the queerness of a certain class of properties. True, I deny that there is, or could be, any such thing as the fact that it might be raining, or the fact that it is possible that it is raining (where again facts are taken to correspond to ways the world might be). But this is not out of any metaphysical doubt that there are such things that might rightly be called epistemic possibilities. On the contrary, I embrace an ontology which includes metaphysically possible worlds,¹⁸ and I embrace the coherence of the idea of a set of such worlds being compatible with a state of knowledge. In that sense, I fully embrace a metaphysics of epistemic possibilities; I am comfortable saying, in a factual tone of voice, that there are possibilities compatible with what I know. The view we

¹⁷ Expressing a state of mind is, remember, distinct from saying that one is in that state of mind.

¹⁸ Indeed my conception of factualism definitionally excludes the possibility of nonfactualism about metaphysical modality in general, for I deploy a 'robust' notion of metaphysical modality in saying what factualism is.

are developing is that, while there are such facts, the language of epistemic possibility does not serve to express them.¹⁹

What of epistemic necessity? Are we nonfactualist about these claims too? It is hard to make a direct case for that view, in the way we have for epistemic possibility. But an indirect case is easy, if we assume what looks plausible, that factual claims are closed under negation. (Factual claims are closed under negation just in case, if a claim ϕ has factualist truth conditions, so does $\neg\phi$.) Since by the duality of epistemic modality, the negation of an epistemic necessity claim is equivalent to an epistemic possibility claim, and epistemic possibility claims are nonfactual, epistemic necessity claims are nonfactual.

5. Question-Sensitivity

The first-order model is significant progress. Given the choice between it and a descriptivist approach, the choice seems to me clear. But I think there is room to improve on the model in an important way. We have one more desideratum to uncover. It arises in connection with a certain problem facing the model.²⁰

Recall Frank, who believes it might be raining in Topeka. Why does he believe this? We could imagine various accounts of how it happened. For instance: He left Topeka this morning and it looked cloudy then. Or the weatherman just now said the chance of rain was 30%. Alternatively, perhaps his evidential situation is more impoverished. Perhaps he has no noteworthy reasons in favor of believing that it's raining in Topeka; rather he merely notices his lack of sufficient reason to believe it isn't raining in Topeka. Perhaps on the way out the door, en route to Topeka, he glances by chance at his umbrella, and the question of rain in Topeka then occurs to him. He realizes he doesn't know whether to expect rain in Topeka. He then comes to think that, well, it might be raining in Topeka.

This last kind of case raises a basic question. What is the difference between Frank's state of mind before the question of rain in Topeka occurs to him and his state of mind after? The question is an uneasy one for the dynamic model. We know, on the model, that Frank's posterior state of belief must be one compatible with the proposition that it's raining in Topeka. But what, we ask, was his prior state of mind? The same: he had no prior beliefs one way or the other as concerns rain in Topeka, so what he believed was compatible with either circumstance. So he has transitioned

¹⁹ At least, not without the help of operators besides epistemic modals. To say *It might be raining* is to say something nonfactual; but to say *For all I know, it might be raining* is to say something factual—it is to describe one's state of knowledge. This is consistent with the view I advocate. (Compare the normative case: *You ought to go to confession* is a normative claim, while *According to Scripture, you ought to go to confession* is not. The meta-ethical expressivist denies the former claim is factual, but makes no such claim about the latter.)

²⁰ A problem along the lines of the one I will describe was noted by Frank Veltman at the University of Michigan Philosophy and Linguistics Workshop of 2006, though I do not know if he would agree with my statement of it. Swanson (2006) also raises a version of this problem.

from its being compatible with his doxastic state that it's raining in Topeka to . . . its being compatible with his doxastic state that it's raining in Topeka. No change. This is wrong: clearly some aspect of Frank's state of mind has changed, and our model ought to capture this change.

We could just as well make the point synchronically, by considering two states of mind at a single time rather than one across time. Compare Frank (in his posterior state) to Hank, a man living across the globe in Rotterdam. Hank has heard of Topeka, and he even knows roughly where it is on the map. But Topeka has no place in his life, and thoughts of Topeka simply have not crossed his mind in years. Like myriad other questions, the question of rain in Topeka today has just not occurred to Hank; and indeed, we stipulate it never will. Does Hank believe it might be raining in Topeka? It seems bizarre to answer affirmatively. It is true, we may stipulate, that for all Hank believes, it is raining in Topeka. For nothing he believes rules that possibility out. But this is just to observe that 'Hank believes it might be raining in Topeka' and 'For all Hank believes, it is raining in Topeka' do not have the same truth conditions. The states of mind of Frank and of Hank, we might say, are alike in as much as for all they each believe, it is raining in Topeka. But they differ in that Frank believes it might be raining in Topeka, whereas that is not so for Hank.²¹

We should like to model the difference. This gives us our last desideratum.

- v. Capture the difference between a proposition's merely being compatible with a state of mind and its being epistemically possible according to that state—that is, possible in the thicker sense connoted by epistemic possibility modals.

But what exactly is that thicker sense? What extra does Frank have over Hank?

It is not easy to say. Should we say that if one believes $\diamond\phi$, one has to be entertaining the proposition ϕ , or to be "seeing the possibility" that this proposition is true? No: we are after a kind of state, not a kind of activity. One does not lose one's beliefs about what might be the case by shifting one's attention to a new topic, or by going to sleep. Should we say instead the requirement is that one has to have entertained the proposition ϕ ? This is a state at least, but it is too demanding. Suppose Jones is looking for his cell phone. He suddenly thinks to himself, *the cell phone might be in the glove compartment of the car*. Thereby a flood of epistemic possibility attributions normally become appropriate for him. For instance:

- Jones thinks the cell phone might be in the car.
- Jones thinks the cell phone might be in a vehicle.
- Jones thinks the cell phone might not be in the living room.

²¹ A third way to put the worry: on the model as so far described, there is no difference between not believing (failing to believe) that a proposition is possible ($\neg B\phi$), and believing that the proposition is not possible ($B\neg\phi$). For believing that a proposition is not possible is just believing its negation ($B\neg\phi$), and this is the only way, on the current model, to fail to believe a proposition is possible. But that is intuitively wrong: my suggestion here is that one can fail to believe a proposition is possible without believing its negation. (This way of stating the problem makes it look analogous to a problem which allegedly afflicts expressivist accounts of normative talk, the so-called "negation problem.")

Jones thinks it might be that one has to go to the car to find the cell phone.

Jones thinks it might be that if one looks in the glove box of the car, one will find the cell phone.

And so on. But clearly we should not say that each of the corresponding propositions here must have been entertained by Jones.

Perhaps we should say this: an agent believes $\Diamond\phi$ just in case ϕ is compatible with his belief state, and moreover ϕ is a nearby consequence of a proposition one has entertained, or recently entertained. Even supposing we can make the idea of a nearby consequence precise, however, we now must worry about amnesiacs who have recently entertained ϕ , but have no recall of that event. If Frank gets hit with a shovel and reverts back to his prior state, he is no longer appropriately described as believing that it might be raining in Topeka, though it is true that he has recently entertained ϕ , and ϕ is compatible with what he believes. This last point suggests the state of believing that something might be the case is not, or not merely, a backward-looking state depending on a history of considered propositions. It has a forward-looking character.

Rather than focusing on states of entertainment or events of seeing possibilities, I want to approach the issue from a different direction. My thermostat is awfully imprecise—it only indicates that the room temperature falls somewhere within a ten degree range. Today it indicates that the temperature in the room is somewhere between 65 and 75 degrees. Thus if you were to ask me if the temperature in the room was 70 degrees (imagine, if you like, you are confirming the accuracy of your newfangled temperature-sensing watch), I could reply: *Well, according to my awfully imprecise thermostat, it might be.* That is to say, my awful thermostat indicates that the temperature in the room might be 70 degrees.

Now like most thermostats, mine does not carry information about the weather in Topeka. The thermostat's occupying any of its information-bearing states in normal conditions is compatible with any state of weather in Topeka. Rain in Topeka, for instance. Thus: for all my thermostat indicates, it's raining in Topeka. Nevertheless it is off to say: *According to my thermostat, it might be raining in Topeka.* (Still worse does it sound to say, in one breath: *According to my thermostat, the temperature in the room might be 70 degrees, and it might be raining in Topeka.*) We have here an asymmetry rather like the one observed between Frank and Hank, but in this case there is little temptation to appeal to what my thermostat may have recently entertained. Rather, in this case it is tempting to press the point that the thermostat is sensitive to some questions, or some subject matters, or a certain class of distinctions, and not others, and that this fact affects the felicity of the relevant uses of the epistemic possibility modal.

Following this line, I suggest that the extra thing that Frank has, and that Hank lacks, is some kind of sensitivity to a question. My thermostat is sensitive to some questions, or some issues, or some distinctions, and not others; so too with Frank and Hank. To count as believing $\Diamond\phi$, ϕ should be compatible with one's beliefs; but in

addition, one's state of belief should also be sensitive to a question for which ϕ is an answer, or partial answer.²² This will be our way of satisfying the fifth desideratum.

What is it to be sensitive to a question in the relevant sense? I think this is the right question to be asking, but I am less sure how to answer it. It seems to be at least this: it is to be equipped with possible states that distinguish possible answers to the question, and it is to be receptive to information which speaks to the question. My thermostat is equipped with possible states that distinguish possible answers to the question, *within what range is the temperature in this room?*, and it is receptive to information which speaks to that question. It is not equipped with possible states that distinguish possible answers to the question, *how is the weather in Topeka?*, and (a fortiori) is not receptive to information which speaks to the question.

This gives us a foothold, but just that. Moving from simple devices back to human beings and their states of belief, matters are less easy to describe. Perhaps if Hank had never heard of Topeka, had no concept of Topeka, and had none of the relations of acquaintance which are prerequisites to having Topeka thoughts, we might comfortably deny that Hank is equipped with possible states that distinguish possible answers to the question. We could say he lacks the resources to frame that question. And surely were that the case, it would indeed be off to say that Hank believes that it might be raining in Topeka. But in our original story, Hank is not so cognitively impoverished. He knows of Topeka. (And about weather.) He is not lacking in conceptual resources. Thus in some sense, he is equipped with possible states that distinguish possible answers to the question of weather in Topeka. Yet our judgment was that the relevant ascription is not correct. Perhaps this is because he is not appropriately equipped in some *other* sense; but I myself can't make this sense out. So I would prefer to press instead the thought that Hank's trouble is that he is not appropriately receptive to information which speaks to the question. The question has not arisen for him, and he would have no interest in it if it did: that itself is sufficient to make an agent insufficiently receptive to information which speaks to the question.

It is sufficient, but let us add that it is not necessary. We can imagine agents who are keenly interested in questions for which they are insufficiently attuned to information about, and so who still may not count as having the relevant epistemic possibility beliefs. Whether something is attuned or receptive to a kind of information is a comparative matter, and judgments about this are sensitive to context. If we are anxiously awaiting learning the result of a cancer test whose result is in a sealed envelope, the salience of the test, whose states are far more sensitive to the relevant question than ours, might trump our claim to be receptive to such information. In that case, although our state of belief is compatible with John's having cancer and we have been

²² The notion of sensitivity I am trying to make precise is distinct from the notion of sensitivity developed by Nozick (1981) pursuant to an analysis of knowledge. On Nozick's use, one's state of belief is sensitive, not to questions, but to propositions: to be sensitive to a proposition in his sense is to be such that one would not believe it, were it not true. (But the two notions are related in at least the following way: if one is sensitive to a proposition in Nozick's sense, this is sufficient, though not necessary, for being sensitive to the question whether the proposition is true in my sense.)

interest in the question, we may prefer to say that we are really not sure whether John might have cancer. (See DeRose 1991.) Likewise, I may hesitate to say that the temperature in my room might be 70 degrees before checking my thermostat. The salience of a test, or device, or a relevant expert may change what counts as being appropriately sensitive to information which speaks to a question.

6. Modeling Question-Sensitivity

Turning from these intuitive considerations back to the modeling question, the plan is to run with the idea that belief is a question-sensitive state. We want to implement this idea formally, in a way that will yield a technical distinction between a proposition's merely being compatible with a state of belief and its being epistemically possible in the more robust sense.

To begin, we need a formal conception of what questions are. I will embrace the conception of questions which goes back to Hamblin (1958), according to which the meaning of a question is taken to be the set of its possible complete answers, where the complete answers to a question are taken to form a set of mutually exclusive propositions.²³ On this conception, a question determines a partition of logical space (or if it is based on some presupposition, a partition of a particular subregion of logical space). The partition gives all the alternative complete answers to the question. The true answer corresponds to the cell of the partition which contains the actual world. Any union of a set of more than one complete answer is an incomplete or partial answer.

Lewis has offered a parallel analysis of *subject matter*. A subject matter is partition of logical space, one that forms equivalence classes of worlds depending on whether the worlds yield the same verdict concerning the subject matter. It divides up logical space, but only so far as concerns those distinctions native to the subject matter. If, for instance, the subject matter is demography, then two worlds demographically alike in all respects will fall into the same cell within the partition determined by the subject matter, though they may differ in any other respect. We can say that a proposition is *about* a subject matter just in case the truth value of the proposition supervenes on the subject matter; equivalently, if the proposition is identical to some unions of cells from the partition. See Lewis (1988a, b) for further details and applications.

²³ In the semantics literature, Hamblin's idea—questions as partitions of a space of options—was developed into its best known form by Groenendijk and Stokhof (1984); see also Groenendijk and Stokhof (1997), Belnap and Steel (1976), Higginbotham (1993, 1996). It is perhaps the leading alternative to the most influential account of questions, namely Karttunen (1977). A central difference between the partition account and Karttunen's account is that on Karttunen's account, it is not assumed that the complete answers to a question must be mutually exclusive. As a result a question might have several true answers. This feature makes Karttunen's treatment of certain phenomena—for instance, so-called 'mention some' questions (e.g. *Where can I find a good sandwich?*) more elegant. The story I am about to tell can, I think, be adjusted to comport with a Karttunenian conception of questions, but I will not be able to go through the exercise here.

As Lewis notes, these two ways of thinking about a partition of logical space—as a question, or as a subject matter—are complementary and come to much the same. I will help myself to both ways of thinking of partitions of logical space. And, at the risk of mixing too many metaphors, I will help myself to a third: I will sometimes call a partition of logical space a *resolution* for logical space. We can think of a resolution as foregrounding some distinctions, bringing them into focus, and backgrounding others. Some propositions will be *visible* at a resolution—that is, about the corresponding subject matter—while the rest will be invisible and off-topic. The cells of a resolution are sets of possible worlds, but we may equally think of the cells themselves as ‘coarse’ possible worlds, worlds that settle some, but not all, questions. I restrict attention to finite resolutions.

The modeling proposal is then this: states of belief are resolution-sensitive. (Or question-sensitive, or subject matter-sensitive.) They are states which are partly relations to some way of resolving logical space, to some way of dividing up the alternatives. Relative to a resolution, a doxastic state will select a set of cells at that resolution as candidates for actuality.²⁴ Formally a doxastic state is, not simply a set of possible worlds, but rather a partial function taking a resolution to a subpartition of that resolution—that is, to a set of coarse possible worlds which is a subset of the resolution, what we may call the *view* of the agent at that resolution. A view gives the doxastically open ‘coarse’ possibilities for the agent at that resolution.

Equivalently, the proposal is that a belief state is a function from questions to answers. The answers may only be partial, eliminating some but not all alternatives. (Or indeed it may eliminate no alternatives. A question is one of its own partial answers—the least helpful one.) And the question reflected by a resolution needn’t be one particular easy to express in language. Depending on how fine it is, it may be preferable to understand it as capturing a family of topically related questions on which the doxastic state takes a stance—as capturing a relatively detailed project of inquiry. We can say a belief state is *sensitive* to a question just in case it is defined on the question (or is defined on a strictly finer question). We take it belief states are not, or at any rate need not be, sensitive to every possible question. A belief state is a partial function, so it may not be total. In realistic cases we may assume it is partial (indeed, defined only on finitely many questions). If a belief state is undefined on a question, say it is *insensitive* to the question. Let us assume that if a belief state is defined on a question, it is defined on any coarsening of the question (where Π' is a *coarsening* of Π just in case every cell of Π' is a union of cells of Π).

A view about a subject matter or question supplies the agent with *commitments* concerning that subject matter or question: these are the propositions true throughout the worlds left open by the agent’s view.²⁵ Within the set of propositions which form the agent’s commitments concerning a subject matter, we can distinguish the visible

²⁴ In the sense that a cell is the *actual cell* just in case it contains the actual world.

²⁵ By which I mean the worlds in the union of the agent’s view.

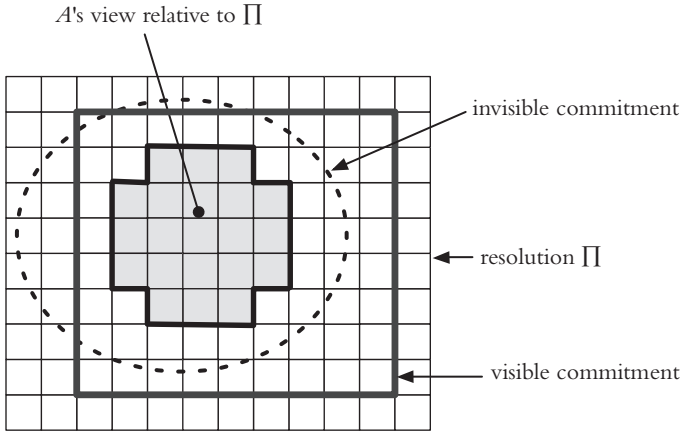


Fig. 10.3. Resolution-sensitive belief.

commitments from the invisible ones. The visible commitments are the ones about the subject matter, in the sense defined above; the invisible ones are the rest. Fig. 10.3 illustrates the distinction between the two sorts of commitments.

Say a proposition is *compatible* with an agent's view (with respect to a resolution his belief state is defined on) just in case it is true at one of the worlds left open by that view. Within the set of propositions which are compatible with an agent's view concerning a subject matter, we can distinguish the visible propositions from the invisible ones. I want to suggest that we use this distinction between ways that

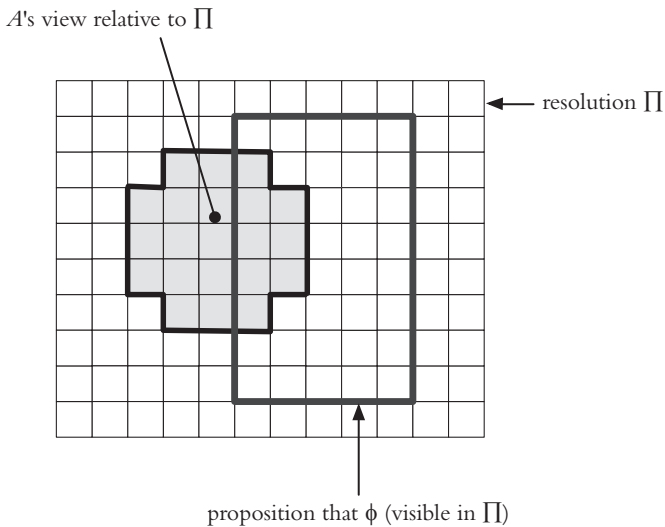


Fig. 10.4. $B_A^\Pi \diamond \phi$: Resolution-sensitive model.

a proposition might be compatible with one's view to model what extra is required to believe that something is possible, in the epistemic sense we are after. To believe that a proposition is possible, or might be, is for the proposition to be compatible with one's view, and moreover for it to be an answer to a question one is sensitive to. (Equivalently, for it to be about a subject matter one is sensitive to; equivalently, for it to be visible at a resolution one is sensitive to.) This kind of state is depicted in Fig. 10.4 below.

We could think of this as a 'pixelated', low-resolution version of the first-order model of epistemic possibility belief. This gives us our fifth desideratum. Our suggestion is that Frank's head looks like this. Hank's state of belief, by contrast, is not defined on a question making the proposition that it's raining in Topeka visible—even though the resolutions his state of belief is defined on yield views entirely compatible with this proposition.

7. Fragmented Inquiry

The question-sensitive model enriches the classic possible worlds model in two ways. First, one has a view of things only relative to a resolution. Second, one has views relative to many distinct resolutions.²⁶

In the latter respect, the resolution-sensitive model can be seen as a more articulated version of the idea of 'fragmented' or 'compartmentalized' belief (Lewis 1982; Stalnaker 1984). On the fragmentation upgrade to the classic possible worlds model, a belief state is still a set of possible worlds, but an agent's total doxastic state is modeled as a set of belief states. One has a belief only relative to a state within this set. Thus belief becomes a three-place relation between a person, a belief state, and a proposition. The key motivation for this upgrade is to avoid the closure of belief under believed material implication and under conjunction. Lewis illustrates the idea nicely:

I used to think that Nassau Street ran roughly east–west; that the railroad nearby ran roughly north–south; and that the two were roughly parallel. . . . Now, what about the blatantly inconsistent conjunction of the three sentences? I say that it was not true according to my beliefs. My system of beliefs was broken into (overlapping) fragments. Different fragments came into action in different situations, and the whole system of beliefs never manifested itself all at once. The first and second sentences in the inconsistent triple belonged to—were true according to—different fragments; the third belonged to both. The inconsistent conjunction of all three did not belong to, was in no way implied by, and was not true according to, any one fragment. That is why it was not true according to my system of beliefs taken as a whole. Once the fragmentation was healed, straightaway my beliefs changed: now I think that Nassau Street and the railroad both run roughly northeast–southwest. (Lewis 1982: 436)

²⁶ Thus the two most recent figures are actually quite incomplete representations of our agent's state of belief: they depict only one resolution, and so just one of the plurality of inquiries that the agent presumably takes a view on.

The basic idea of fragmentation is, I think, right. That is why the resolution-sensitive model encodes a version of the idea: understand the ‘fragments’ to correspond to the various distinct views associated with the various distinct questions one’s state of belief is defined on. If situations for action can be associated with different questions or subject matters, we can say that the aspect of an agent’s state of belief relevant to explaining their behavior in a particular situation is determined by the question they associate with that situation—with the way in which they carve up the options. Since the views of an agent relative to one way of carving up the options may not line up with his views relative to another way of carving up the options, an agent may have inconsistent views—but without thereby counting as believing everything (as is the case on the classic possible worlds model when one has inconsistent beliefs).

In general, we want our views concerning our disparate inquiries to cohere and be compatible. When we see incompatibilities, we shift our views so that they cohere. Fragmentation occurs because it is a non-trivial matter to bring our inquiries together into a single state of mind.

Why is this a non-trivial matter? If we want, we can use the resolution-sensitive model to formulate an answer to this question. Wherever you have a finite space of alternatives of the sort presented by a resolution, you can give a measure of how much information, in bits, would be needed to reduce those alternatives to one. This is just the logarithm, to the base 2, of the number of alternatives. Call this number in bits the *information potential* of a resolution. Now although the idea of resolution-sensitivity does not come with any specific commitments about the detailed form by which the content of belief is represented, we might wish to construe it so that it does at least demand that the representational vehicle of belief at least have the complexity to encode, in bits, the information potentials of the various resolutions the agent is sensitive to. Resolutions impose what we could call an *encoding cost*. This is a way, albeit a highly abstract way, that this framework for representing belief imposes a constraint on a model of the mechanism of representation. Then we can say that fragmentation happens because bringing out disparate inquiries together into a single state of mind carries an encoding cost we cannot afford. It would require a state of mind at a higher resolution than we are capable of.

So the resolution-sensitive model has the advantage of the fragmentation model, *vis-à-vis* the closure properties that motivate fragmentation. One’s commitments are not generally closed under conjunction, because we cannot speak of “one’s commitments” full stop. Rather, various questions give rise to various views, which give rise to various packages of commitments. Relative to a particular question, an agent’s commitments will be closed under conjunction; but not necessarily so, for commitments stemming from different questions.

And the resolution-sensitive model arguably has a further advantage. The problem of the closure of belief under conjunction, the problem eased by fragmentation, is not the most problematic closure property afflicting the classic possible worlds account. The heart of the problem of logical omniscience results from the closure of belief

under truth-conditional entailment. The simple fragmentation model does nothing to ease this problem. But the resolution-sensitive model has resources here. Relative to a subject matter, one's commitments are indeed closed under entailment; but one's visible commitments are not. This affords us the option, then, of using the distinction between visible and invisible commitments to articulate a sense in which realistic agents are not logically omniscient.

Let me illustrate a simple application in this connection, using an example from Stalnaker (1984). Suppose the following is true:

- (13) William III of England believed, in 1700, that England could avoid war with France.

The proposition that England could avoid war with France truth-conditionally entails the proposition that England could avoid nuclear war with France. Thus it appears that on the classic possible worlds model, it follows from (13) that

- (14) William III of England believed, in 1700, that England could avoid nuclear war with France.

But this seems absurd. Perhaps in some sense William III believed something which committed him to the truth of a proposition concerning nuclear war, but the belief ascription (14) sounds wrong. We can offer the following explanation on the resolution-sensitive upgrade: William III's state of belief is not sensitive to questions concerning nuclear war. Equivalently, it is not sensitive to subject matters about nuclear war. The proposition that England could avoid nuclear war with France was nowhere among his visible commitments. The belief ascription suggests otherwise, and this is why it is defective.

8. Subject Matters in Language

Except, we need to explain in virtue of what the belief ascription 'suggests otherwise'. When we fragment belief, whether in the style of Lewis and Stalnaker or in the resolution-sensitive manner I have advocated, we incur a semantic obligation. Belief ascriptions appear to relate to individuals and propositions, but advocates of fragmentation are committed to the idea that the underlying reality involves a further relativity to a belief state (Lewis, Stalnaker) or to a question (Yalcin). To my knowledge, neither Lewis nor Stalnaker has attempted to pay this debt. But to fully vindicate our treatment of (14), we require a semantics for belief reports in the resolution-sensitive setting. In this section I mainly wish to acknowledge this debt, and make a small down payment.

There are various ways one might attempt to connect the technical model we have described to a semantics for belief reports, and in particular reports embedding possibility modals. Without being able to motivate it fully, let me briefly sketch one approach. It uses the resources of alternative semantics, an approach developed in connection

with the semantics of questions and of focus (Hamblin (1973), Rooth (1985, 1992), among others).

An alternative semantics recursively defines two functions. First there is the usual *semantic value function* $\llbracket \cdot \rrbracket$, which takes expressions to extensions (relative perhaps to some points of evaluation). Second there is what I will call the *alternative semantic value function* $[\cdot]$, which maps an expression to a set of alternatives appropriate to the type of the expression. It is easiest to illustrate the concept with an example from the focus literature. If our sentence is:

(15) Mary likes *Sue*.

with ‘Sue’ focused, then the alternative semantic value of the sentence is a certain set of propositions:

$$\{\text{Likes}(\text{Mary}, \gamma) : \gamma \in E\}$$

(Where E is the domain of individuals). Intuitively, these propositions constitute the alternatives to Mary liking *Sue*. Roughly speaking, they are the various propositions which would be determined by the replacement of ‘Sue’ in the sentence with a name for another individual.²⁷ This set includes (e.g.) the proposition that Mary likes Alice and the proposition that Mary likes John, but it does not include (e.g.) the proposition that Sue likes Mary. Note this set also includes the proposition that Mary likes Sue: we assume the proposition determined by the ordinary semantic value of the sentence is an element of its alternative semantic value. The alternative semantic value of a sentence is compositionally determined from the alternative semantic values (and regular semantic values) of its components, but we remain agnostic on those details. (See Rooth (1985), Kratzer (1991a) for proposals and further references.)

The alternative semantic value of a sentence is a partition of logical space. So it is a resolution, or subject matter, or question. I wish to exploit this fact in the semantics of ‘believes’. The alternative semantic value of the complement of ‘believes’ supplies a question. By the resolution-sensitive model, the subject of a belief ascription, together with the world of evaluation, determines a belief function—a function from questions to answers, or from subject matters to views. Semantically then, we can say that in evaluating belief ascriptions, we evaluate the agent’s belief function relative to the question or subject matter equivalent to the alternative semantic value of the complement clause.

So consider an agent A , whose belief function is B relative to w . If Π is a partition this function is defined on, then $B(\Pi)$ delivers the agent’s view relative to Π . On the model we have developed $B(\Pi)$ is a subpartition of Π —a set of cells from Π —but within the semantics, it will be simpler to use a function that directly gives the union of the relevant subpartition. So take it B is a function from a partition of logical space

²⁷ This is rough because an individual x does not need to have a name in the language in order for the proposition that Mary likes x to be among the alternatives to (15).

to a set of worlds, a set equivalent to a union of cells from that partition. Then we can give a possible worlds semantics for the belief operator as follows:

$$\llbracket B_A \phi \rrbracket^w = 1 \text{ iff } \forall w' \in B_A^w(\llbracket \phi \rrbracket) : \llbracket \phi \rrbracket^{w'} = 1$$

In respect of its universal quantification over worlds, this parallels the classic account going back to Hintikka (1962). Note this semantics encodes the requirement that the proposition expressed by ϕ be relevantly visible for the agent, because B_A^w must be defined as a question for which ϕ is an answer in order for the ascription to be true. The ascription (14) would fail this requirement, because William III's belief state is not defined on any partition that could correspond to [**England can avoid nuclear war with France**].

An immediate problem with this semantics is that it does not easily mesh with the first-order picture of epistemic possibility belief we have been defending. It is wedded to the notion that the belief operator must combine with something determining a condition on possible worlds, but we have seen reason to doubt that idea. Fortunately, it is not hard to adapt the semantics so that it expresses what we want. The following style of adjustment is motivated by Yalcin (2007) (see also MacFarlane (this volume); Kolodny and MacFarlane (2010)). Let semantic values be defined relative, not just to worlds, but also to *states of information* s , which we will take to be sets of possibilities (construing possibilities now as partition cells). For most sentences this new parameter will be idle, but it comes into action in connection with the semantics of attitudes and epistemic modals. Let epistemic possibility clauses existentially quantify over the possibilities open according to the information state parameter, as follows:

$$\llbracket \Diamond \phi \rrbracket^{w,s} = 1 \text{ iff } \exists w' \in s : \llbracket \phi \rrbracket^{w',s} = 1$$

Next, adjust the semantics for belief operators so that, in addition to effecting universal quantification over a certain view as above, these operators shift the value of the information parameter to that view. As follows:

$$\llbracket B_A \phi \rrbracket^{w,s} = 1 \text{ iff } \forall w' \in B_A^w(\llbracket \phi \rrbracket) : \llbracket \phi \rrbracket^{w',B_A^w(\llbracket \phi \rrbracket)} = 1$$

Last, take it that the alternatives to $\Diamond \phi$ reduce to the alternatives to ϕ :

$$\llbracket \Diamond \phi \rrbracket = \llbracket \phi \rrbracket$$

With all this, it follows that:

$$\llbracket B_A \Diamond \phi \rrbracket^{w,s} = 1 \text{ iff } \exists w' \in B_A^w(\llbracket \phi \rrbracket) : \llbracket \phi \rrbracket^{w',B_A^w(\llbracket \phi \rrbracket)} = 1$$

which is just what we want. Epistemic possibility beliefs come out as first-order states of mind, and one is in this state of mind only if one is sensitive to a question for which the relevant proposition is an answer.²⁸

²⁸ It may strike some readers as questionable to assume that the alternatives to $\Diamond \phi$ reduce to the alternatives to ϕ . Perhaps the facts about focus, for instance, will teach us otherwise. That is certainly possible. But, first, theoretically we can take $\llbracket \]$ to be a *sui generis* function for use in modeling the question-sensitivity of thought, a function related, but not identical to, the alternative semantic values relevant to focus interpretation. Second, we really only need the weaker assumption that $\{w : \llbracket \phi \rrbracket^{w,\emptyset} = 1\} \in \llbracket \Diamond \phi \rrbracket$ to get what

A second, rather different approach to giving resolution-sensitive semantics for belief reports would be to treat ‘believes’ as expressing a three-place relation between a person, a proposition, and a resolution, but where the latter component is not supplied by any secondary kind of semantic value. Since (focus structure aside) there is no obvious linguistic constituent corresponding to a resolution in an ordinary belief report, this view would take the value of resolution variable to be either indexically supplied by context, or (more plausibly) to be tacitly existentially quantified over. So the relevant logical form would be something like:

$$\exists \Pi : B(A, \Pi, p)$$

So-called ‘hidden-indexical’ analyses of belief reports give semantics in a structurally parallel fashion. Traditionally that approach has been motivated by a desire to reconcile a Russellian view of propositions with Fregean intuitions about belief reports. Belief is taken to be a three-place relation, but the third argument place is occupied by a mode of presentation (concept, Fregean sense), not a question or subject matter. There is an extensive literature on this approach, as it has played a non-trivial role both in debates about Frege’s puzzle and in debates about the extent to which syntax constrains semantics. (For relevant discussion see Schiffer (1977, 1992, 2003), Crimmins and Perry (1989), Crimmins (1992), Ludlow (1996), Stanley (2000), Recanati (2002), Hall (2008).)

I see some room for the hidden-indexical and resolution-sensitive approaches to dovetail. There is a perfectly intelligible sense in which a resolution supplies a mode of presentation of a proposition in the possible worlds setting. A proposition is a rule for dividing maximally specific possibilities, but a resolution reflects what features of those possibilities are at issue, foregrounding the distinctions that matter and backgrounding the distinctions that do not. A single proposition may be ‘presented’ relative to various partitions of logical space, and these will highlight different features of the kinds of possibilities the proposition rules in and rules out.

But if you enjoy thinking of subject matters as senses, realize you can do this already on the alternative semantics given above. Second, I doubt resolutions alone would be of great use in solving Frege’s puzzle, the chief motivation for classical hidden-indexical theories. (I myself would prefer a solution to Frege’s puzzle that adverts to unstructured sets of certain finer-grained possibilities—trading in possible worlds for possible world, sequence-of-individual pairs—but that is another story. See Ninan (2008) and Cumming (2008) for views in the spirit of the approach I favor.) Third, if the relevant alternative semantic values are independently motivated by the semantics of questions and of focus, as I suspect, an alternative semantics is more elegant. It uses structure already needed elsewhere, and does not posit any tacit quantifiers whose scoping possibilities we would have to artificially restrict. Only if

we want. And indeed, it seems not implausible that one of the alternatives to ‘Bob *might* be in his office’ is ‘Bob *is* in his office’.

an alternative semantics for belief reports cannot be made out should we turn, then, to a hidden-indexical analysis.

9. Credal Expressivism

I advertised a defense of nonfactualism about, not just pure claims of epistemic modality—*might* and *must* claims—but also graded claims of epistemic modality—of *probably* claims. But my theory-building so far has been obsessed with epistemic possibility.

Probability operators give rise to epistemic contradictions; they give rise to puzzles of agreement and disagreement; they give rise to the tensions of eavesdropping; and they bear straightforward logical relations to the epistemic modals *might* and *must*. It is also not at all hard to motivate the idea that believing something is probable is a first-order state of mind. Indeed, that is arguably the default view in the Bayesian literature. If you ask a Bayesian to tell you what sort of constraint a credence function C has to satisfy in order for an agent to count as believing a proposition p is probable, I wager she will say the constraint is this: $C(p) > .5$. (Or at any rate, she will say $C(p)$ must take some highish value.) This would be to say that to believe a proposition is probable is to be in a doxastic state of mind modelable by a probability space, one whose measure assigns that proposition a relevantly high value. And that is already to say that this state of mind is not fundamentally second-order in character. It is not a matter of one's credence in a proposition *about* one's credence; neither is it a matter of one's credence in some proposition about one's evidence, as the descriptivist proposals discussed above would most naturally recommend. It is simply a matter of how one is credally related to p .²⁹

So the signs point to an expressivist, nonfactualist view about probability claims, at least if we have been right so far about epistemic possibility. How exactly to extend the resolution-sensitive model to make room for probabilities is a subtle matter, and I regret I must save detailed discussion for elsewhere.³⁰ For now at least, we can briefly mention some of the obvious moves to make, and some obvious contours of the view to be spelled out. Lay down probability measures over the resolutions the agent's state of mind is defined on, measures which place all the probability mass within the boundaries of the relevant views. Just as an agent's acceptance of an epistemic possibility claim is a function of whether a visible proposition is compatible with the relevant view of the agent, his acceptance of an epistemic probability claim will be a function of whether a visible proposition gets a high enough probability according to the probability measure of the relevant view. One expresses one's highish credence in a proposition when one says something is likely. But thereby, one does not say that one

²⁹ Thus the Bayesian picture itself needn't be tied to anything like a descriptivist account of probability talk, as is sometimes informally assumed (e.g. by Jeffrey (2004); see n. 3 above).

³⁰ See Yalcin (2007, 2010) for some relevant discussion, and Swanson (2006) for a related approach.

is in a state of high credence. One does not describe oneself, or one's evidence. One expresses a property of one's state of mind—not a proposition, but a condition on a probability space—with the intention of coordinating one's interlocutors' states on the satisfaction of that property. In a given context, there may be intelligible questions of rationality or advisability to be raised about such claims. But the question of truth as it arises for ordinary factual informational content does not arise here, any more than it would for epistemic possibility claims.

10. Nonfactual Parameters

In meta-ethics, the name 'expressivism' is sometimes attached to a cluster of theses to the effect that normative claims are fancy riffs on 'Boo!' and 'Yay!'—that they are essentially yelps in linguistic dress, primarily 'expressing' some non-contentful attitudes pro and con and having no compositional semantics. I find that view—perhaps better called *emotivism*—totally unbelievable, and I hope it is clear that the theory I have defended has hardly anything in common with it. On the contrary, the view I have defend is partly *motivated* by considerations from compositional semantics.

But some readers will find my position puzzling. I have said that epistemic possibility claims are nonfactual, and yet I offer a semantics for the epistemic possibility modal which looks like this:

$$\llbracket \Diamond \phi \rrbracket^{w,s} = 1 \text{ iff } \exists w' \in s : \llbracket \phi \rrbracket^{w',s} = 1$$

And this seems awfully like truth-conditional semantics. Isn't '1' just another name for The True, after all? Am I then not really a factualist, albeit one with some overblown expressivist rhetoric? Or worse—since the recursive definition of '1' above involves relativity to a state of information—am I not really a *relativist* about epistemic modality?

The distinction between relativism and expressivism has been questioned (Field 2009) but there is, I think, a way to make it out. The crucial point is to recognize a distinction between a view about compositional semantic value and a view about informational content. The notion of *truth at a point of evaluation* is a technical notion from semantics. We use it to articulate our tacit semantic competence, a competence usefully characterized, at the level of sentences, by an ability to effect a distinction between a certain space of points. The structure of these points depends on the contingencies of human language—for instance, on what operators the language contains. Separate from this technical notion of truth is the notion of truth as it applies to propositions, or items of informational content. Insofar as there is a concept of truth relevant to whether a claim or state of mind is factual in character, it is the latter notion we should be concerned with.

I said above that I follow Lewis (1980) in taking it that where semantics meets pragmatics, we need not assume that the semantic value of a declarative sentence is its informational content. We need only assume that the informational content of a

sentence is recoverable from its semantic value, together with features of context and whatever standing pragmatic knowledge there may be. Now on one familiar way of interpreting the two-dimensional semantics of Kaplan (1989), there is a simple recipe for recovering the informational content of a sentence from its semantic value. If the semantics of ϕ is the two-dimensional intension $\llbracket \phi \rrbracket^{c,i}$, where c is a context and i is index (a tuple of features of context shiftable by operators in the language), then (on one view) the informational content of the sentence in context is:

$$\lambda c. \llbracket \phi \rrbracket^{c,i_c} = 1$$

(where i_c is the index determined by the context). This is the *diagonal proposition* determined by the sentence. This notion of informational content is naturally paired with a notion of truth distinct from, but definable in terms of, truth at a point of evaluation:

$$\phi \text{ is true at a context } c \text{ iff } \llbracket \phi \rrbracket^{c,i_c} = 1$$

Whether a sentence is *true at a context* tracks exactly when the diagonal proposition it expresses is true. Diagonal content is a kind of factual content: it discriminates between metaphysically possible ways a context might be.

Now the point to be clear about is that if one holds that the diagonal corresponds to the informational content of the sentence, then one must already acknowledge a gap between truth at a point of evaluation, the notion appropriate to compositional semantics, and the notion of truth appropriate to content. And indeed, it should be clear that the same is true if one holds that the informational content of the sentence is the horizontal proposition of the sentence in context. Thus the idea that there is a gap between truth at a point of evaluation and truth as it applies to the objects playing the role of content is not a new idea. It is built into the usual ways of interpreting Kaplan's framework.

Expressivism about a domain of discourse is sometimes characterized as the view that the claims of the discourse are not truth-apt. But once we recognize this distinction between two notions of truth—the notion of truth belonging to formal semantics and the notion of truth belonging to the theory of content, or as we might have it in a two-dimensional setting, *truth at a point of evaluation* and *truth at a context*—the expressivist position can be more subtle than this. The expressivist can give a recursive definition of truth at a point of evaluation in the usual way for his discourse, thereby preserving compositionality, but he can reject the demand to give truth conditions in the more robust sense, the one appropriate to ordinary factual informational content. To illustrate more precisely: if our starting point is two-dimensional semantics, and the informational content of a sentence is taken to characteristically be its diagonal, what the expressivist can deny is that the notion of truth at a context is well-defined for the sentences of his target discourse. This is to say that the claims of the discourse do not have informational content—or at least, that they do not have it in the ordinary, factual sense, the sense which a well-defined diagonal makes out.

Bringing this to bear on the above semantics for epistemic possibility, the claim would be that according to the expressivist,

$$\lambda c. \llbracket \Diamond \phi \rrbracket^{c, w_c, s_c} = 1$$

is not well-defined, because there is no such thing as s_c —no such thing as ‘the information state of the context’. While of course there presumably are information states in the context (the states of the interlocutors, for instance), we stipulate as theorists that the role of this parameter is not to represent any of them. Rather, it is a *non-factual parameter*. Unlike the world coordinate or time coordinate (if such there be), this parameter does not correspond to any possible feature of context. There is no ‘initializing’ it; rather, what is communicated is variable with respect to its value. The expressivist can make this move, compatible with telling a story that makes communication employing the claims of the discourse intelligible. As we have said: relative to context, an epistemic possibility claim determines a condition, or property, on states of information—on states of mind. It is the satisfaction of this property that the speaker aims to coordinate his listeners on. The speaker thereby expresses a feature of his state of mind, and does so without describing himself, or the world.

This approach is, I think, relevantly analogous to that of Gibbard (1986; 1990; 2003; see especially 2003: ch. 6). Gibbard has his own nonfactual parameter—a parameter for a system of norms, or in more recent work, for a maximally complete plan for action, or *hyperplan*. One’s state of belief is plan-laden, representable as a set of possibility-hyperplan pairs. Normative sentences serve chiefly to divide the space of hyperplans. Crucially for Gibbard, when hyperplans appear in the formal semantics, they are not taken to correspond to any feature of context. The truth of a normative claim is not, for instance, a matter of what norms the participants in the conversation accept. If it were, such claims would then merely describe the norms endorsed by the interlocutors, and that is not the idea. Rather, the parameter is nonfactual. The question of which hyperplan is the *right* one to use in evaluating a sentence is not one determined by the facts of the context. Rather, on Gibbard’s view, it is a practical question, a question about what to do. Gibbard uses hyperplan parameter to say what it is we are up to in communicating with normative sentences:

Plan-laden conviction, though, pertains not to the diagonal of the character matrix, but to the diagonal with a hyperplan dimension added. It is given by a diagonal plane in the extended character matrix. We can talk now of the *extended import* of a thought or other concept. When I accept what someone says on his authority, it is the extended import of what he says that gets communicated; I come to accept it. (Gibbard 2003: 132)

On Gibbard’s picture, purely factual content (purely factual import) corresponds essentially to diagonal content of the sort already described. Even if, for the sake of uniformity, all content is represented in terms of sets of possibility-hyperplan pairs—*extended imports* in Gibbard’s terminology—we can distinguish factual from nonfactual content, for content which is purely factual does not vary with the choice

of hyperplan. That means there is a perfectly interesting, factually-oriented, well-defined notion of *truth at a context* applicable to all factual discourse in Gibbard's setting—basically, the Kaplanian notion.³¹ This is just as we would want, I presume.

But we have no such definition available for normative thought and talk, if Gibbard is right. If the extension of a sentence varies with choice of hyperplan, it is normative. Since context does not determine a choice of hyperplan, we cannot define a diagonal from the semantics of the sentence without fundamentally distorting the import of the claim. Instead of the factulist truth conditions that a diagonal would supply, what we have is an object—a non-trivial extended import—which is variable with respect to hyperplan. It is this object Gibbard uses to say what agents engaging in normative talk are communicating and coordinating on.

Expressivists have felt cornered on the question of truth. If they deny that normative discourse is truth-apt, they fall prey to the Frege-Geach problem; if they affirm it is, it becomes hard to see how their from that of the factulist. If they go minimalist about truth, their position applies to all of language, not just a fragment of it. I am suggesting that the expressivist—or one attractive brand of expressivist, anyway—can say what is distinctive about his position by exploiting the independently motivated distinction between compositional semantic value and informational content. And I have suggested that one concrete way of doing this begins from the perspective of two-dimensional semantics. Like the factulist, the expressivist defines a notion of truth at a point of evaluation. Thereby he vindicates the undeniable compositionality of natural language. But unlike the factulist, he rejects the view that the sentences of the relevant discourse are apt for truth in a richer sense, the sense of truth which applies to factual information content—the kind of content whose main business is to rule out ways things might be.

References

- Beaver, David (2001). *Presupposition and Assertion in Dynamic Semantics*. Studies in Logic, Language and Information. CSLI Publications.
- Belnap, Nuel, and Steel, Thomas (1976). *The Logic of Questions and Answers* (New Haven: Yale University Press).
- Crimmins, Mark (1992). *Talk about Beliefs* (Cambridge, Mass.: MIT Press).
- and Perry, John (1989). “The prince and the phone booth: Reporting puzzling beliefs”, *Journal of Philosophy* 86(12): 685–711.
- Cumming, Samuel (2008). “Variabilism”, *Philosophical Review* 117(4): 525–54.
- DeRose, Keith (1991). “Epistemic possibilities”, *Philosophical Review* 100(4): 581–605.
- Dretske, Fred (1981). *Knowledge and the Flow of Information* (Cambridge, Mass.: MIT Press).
- Egan, Andy (2007). “Epistemic modals, relativism and assertion”, *Philosophical Studies* 133(1): 1–22.

³¹ One precise version of such a definition: ϕ is true at c just in case $\llbracket \phi \rrbracket^{c,c,h} = 1$, for any h (where h is a variable over hyperplans).

- Hawthorne, John, and Weatherson, Brian (2005). “Epistemic modals in context”, in G. Preyer and G. Peter (eds.) *Contextualism in Philosophy* (Oxford: Oxford University Press) 131–68.
- Evans, Gareth (1983). *The Varieties of Reference* (Oxford: Oxford University Press).
- Field, Hartry (2009). “Epistemology without metaphysics”, *Philosophical Studies* 143(2): 249–90.
- von Fintel, Kai, and Gillies, Anthony S. (2008). “CIA leaks”, *Philosophical Review* 117(1): 77–98.
- Gibbard, Allan (1986). “An expressivistic theory of normative discourse”, *Ethics* 96(3):472–85.
- (1990). *Wise Choices, Apt Feelings* (Cambridge, Mass.: Harvard University Press).
- (2003). *Thinking How to Live* (Cambridge, Mass.: Harvard University Press).
- Gillies, Anthony (2004). “Epistemic conditionals and conditional epistemics”, *Noûs* 38: 585–616.
- Groenendijk, Jeroen, and Stokhof, Martin, (1984). *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD thesis, University of Amsterdam.
- (1997). “Questions”, in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language* (Amsterdam: Elsevier Science), 1055–124.
- Hall, Alison (2008). “Free enrichment or hidden indexicals”, *Mind and Language* 23(4): 426–56.
- Hamblin, Charles L. (1958). “Questions”, *Australasian Journal of Philosophy* 36(3): 159–68.
- (1973). “Questions in Montague English”, *Foundations of Language* 10(1): 41–53.
- Heim, Irene (1992). “Presupposition projection and the semantics of attitude verbs”, *Journal of Semantics* 9: 183–221.
- Higginbotham, James (1993). “Interrogatives”, in K. Hale and S. J. Keyser (eds.), *The View from Building 20* (Cambridge, Mass.: MIT Press), 195–227.
- (1996). “The semantics of questions”, *The Handbook of Contemporary Semantic Theory* (Oxford: Blackwell), 361–83.
- Hintikka, Jaakko (1962). *Knowledge and Belief: An Introduction to the Logic of The Two Notions* (Ithaca, NY: Cornell University Press).
- Jeffrey, Richard (2004). *Subjective Probability: The Real Thing* (Cambridge: Cambridge University Press).
- Kaplan, David (1989). “Demonstratives”, in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press), 481–563.
- Karttunen, Lauri (1977). “Syntax and semantics of questions”, *Linguistics and Philosophy* 1: 607–53.
- Kolodny, Niko, and MacFarlane, John (2010). “Ifs and oughts”, *Journal of Philosophy* 107(3): 115–43.
- Kratzer, Angelika (1977). “What ‘must’ and ‘can’ must and can mean”, *Linguistics and Philosophy* 1: 337–55.
- (1981). “The notional category of modality”, in H.-J. Eikmeyer and H. Rieser (eds.), *Words, Worlds, and Contexts: New Approaches to Word Semantics* (Berlin: de Gruyter), 38–74.
- (1991a). “The representation of focus”, *Semantics: An International Handbook of Contemporary Research* (Berlin: de Gruyter), 825–34.
- (1991b). “Modality”, in A. von Stechow and D. Wunderlich (eds.), *Semantics: An International Handbook of Contemporary Research* (Berlin: de Gruyter), 639–50.
- Lewis, David K. (1979). “Scorekeeping in a language game”, *Journal of Philosophical Logic* 8: 339–59.

- Lewis, David K. (1980). "Index, context, and content", in S. Kanger and S. Ohman (eds.), *Philosophy and Grammar* (Reidel), 79–100.
- (1982). "Logic for equivocators", *Noûs* 16(3): 431–41.
- (1988a). "Statements partly about observation", *Philosophical Papers* 17(1): 1–31.
- (1988b). "Relevant implication", *Theoria* 54: 161–74.
- Ludlow, Peter (1996). "The adicity of 'believes' and the hidden indexical theory", *Analysis* 56 (2): 97–101.
- MacFarlane, John (this volume) "Epistemic modals are assessment-sensitive", Ch. 5.
- Moore, G. E. (1962). *Commonplace Book 1919–1953* (London: George Allen, and Unwin).
- Ninan, Dilip (2008). *Imagination, Content, and the Self*, PhD thesis, Massachusetts Institute of Technology.
- Nozick, Robert (1981). *Philosophical Explanations* (Cambridge, Mass.: Belknap Press).
- Price, Huw (1983). "Does 'probably' modify sense?", *Australasian Journal of Philosophy* 61(4): 396–408.
- Recanati, Francois (2002). "Unarticulated constituents", *Linguistics and Philosophy* 25(3): 299–345.
- Rooth, Mats E. (1985). *Association with focus* (Cambridge, Mass.: University of Massachusetts).
- (1992). "A theory of focus interpretation", *Natural language semantics* 1(1):75–116.
- Schiffer, Stephen (1977). "Naming and knowing", *Midwest Studies in Philosophy* 2(1): 28–41.
- (1992). "Belief ascription", *Journal of Philosophy* 89(10): 499–521.
- (2003). *The Things We Mean* (Oxford: Oxford University Press).
- Stalnaker, Robert (1984). *Inquiry* (Cambridge: MIT Press).
- (1970). "Pragmatics", *Synthese* 22: 272–89.
- Stanley, Jason (2000). "Context and logical form", *Linguistics and Philosophy* 23(4): 391–434.
- (2005). "Fallibilism and concessive knowledge attributions", *Analysis* 65(2): 126–31.
- Swanson, Eric (2006). *Interactions with Context*, PhD thesis, Massachusetts Institute of Technology.
- Veltman, Frank (1985). *Logics for Conditionals*, PhD thesis, University of Amsterdam.
- (1986). "Data semantics and the pragmatics of indicative conditionals", in E. Traugott, A. ter Meulen, J. Reilly, and C. Ferguson (eds.), *On Conditionals* (Cambridge: Cambridge University Press), 147–68.
- (1996). "Defaults in update semantics", *Journal of Philosophical Logic* 25(3): 221–61.
- Yalcin, Seth (2007). "Epistemic modals", *Mind* 116 (464): 983–1026.
- (2008). *Modality and Inquiry*, PhD thesis, Massachusetts Institute of Technology.
- (2010). "Probability operators", *Philosophy Compass* 5(11): 916–37.
- and Joshua Knobe. (2010). "Fat Tony Might Be Dead. An Experimental Note On Epistemic Modals." Unpublished, <http://www.semanticsarchive.net/Archive/TdiZjA3N/>