

8. Vagueness and Pragmatics

As we have noted several times, if Louis is a penumbral case of baldness, then many competent speakers will not be disposed to assent to any of (1) through (3), though they will assent to (4).

- (1) Louis is bald.
- (2) Louis is not bald.
- (3) Louis is bald or Louis is not bald.
- (4) It is not the case that Louis is bald and that he is not bald.

There is some complication here when (4) is used idiomatically, but once we allow for that usage, these four reactions provide a fairly basic piece of data that a good theory of vagueness ought explain. Most theorists have something like explanations of our reactions to (1) and (2). Some are built to explain our reactions to (3) – theories that advocate reforming classical logic to accommodate data concerning vagueness are paradigm cases of this. Some are built to explain our reactions to (4) – theories that stress penumbral connections, like supervaluationism and epistemicism are paradigm cases of this. What is trickier is to provide an explanation of our reactions to both (3) and (4). In this chapter, I provide such an explanation via a broader theory of the pragmatics of compound sentences. The existence of this explanation will provide some motivation (and cover) for the semantic theory of vagueness outlined in the next chapter. If all of the anomalies can be given a pragmatic explanation, then there is no reason to adjust our logic to meet them. I will also argue that only a semantic theory of a broadly supervaluational nature could work smoothly with this pragmatic theory, but that is more contentious, and less important. The central point is that we can explain the data behind the anomalies without resort to logical deviance.

The core idea will be a development of one outlined by Kit Fine (1975: 140) and Rosanna Keefe (2000: 164). Fine and Keefe are both supervaluationists, but the theory they present *looks like* it could work independently of the supervaluational framework. Louis's being bald is not a sufficient condition for you to properly assert he is bald - rather he must be *determinately* bald. Or, perhaps more perspicuously, it must be determinately true that he is bald. For Fine and Keefe, being determinately true means being supertrue, but may mean something different to you if, perchance, you disavow supervaluationism. However we understand determinacy, we should agree that a simple sentence like (1) is assertable only if it is supertrue. Assuming other factors are equal, the audience is interested in the state of Louis's hair, you have adequate epistemic access to that state, and so on, (1)'s being determinately true will also be a sufficient condition for its proper assertion. We will assume from now on that those conditions are met. We will write *Determinately S* as $\Box S$ throughout in what follows, noting where necessary what assumptions we are making about its logic.

Now comes the crucial step. If that was all you were told, you would think that for a disjunction $\Box(A \text{ or } B)$ could be properly asserted iff it were determinately true, just like all other sentences. But, Fine and Keefe suggest, perhaps we take the condition in which it can be properly asserted to be different to this. We think (rightly or wrongly) that it can only be properly asserted if $\Box A \vee \Box B$ is true, and not merely if $\Box(A \vee B)$ is true. The bulk of this chapter consists of a development of this idea, and a defence of that development.

According to the theory presented thus far, there is a fairly mechanical procedure for connecting simple sentences with their assertion conditions. The suggestion, then, is that we work out the assertion conditions for compound sentences by applying that procedure not to whole sentences, but to their parts. This is how we get the assertion condition for $\Box(A \text{ or } B)$ being $\Box A \vee \Box B$. To which

parts should we apply this procedure? Well, perhaps there are no hard and fast rules about this. Perhaps context determines whether the procedure should be applied to a whole sentence or to its sentential parts. You might expect that this will mean context determines whether sentences like (3) and (4) can be properly asserted. This is right in the case of (3). If its assertion condition is (3a) then it can be asserted, if it is (3b) then it cannot. ((3a) is the case where the procedure is applied to the whole sentence, (3b) where it is applied to the parts.)

- (3a) \square (Louis is bald or Louis is not bald)
 (3b) \square (Louis is bald) or \square (Louis is not bald)

However, this speculation would be wrong about (4). No matter how or where we apply the mechanical procedure, the assertion condition for (4) that is generated is true, as (4a) to (4c) illustrate.

- (4a) \square (It is not the case that Louis is bald and that he is not bald)
 (4b) It is not the case that \square (Louis is bald and that he is not bald)
 (4c) It is not the case that \square (Louis is bald) and that \square (he is not bald)

Apply our little procedure to (4) any way you like, and provided you've started with a broadly classical theory like supervaluationism or epistemicism, you will predict that (4) can be properly asserted. So the theory sketched by Fine and Keefe looks like it has a chance of capturing some rather interesting data.

The two core aims of this paper are to show that Fine and Keefe's theory, as amended and extended, can (a) explain all the data about our reactions to compound sentences involving vague clauses like *Louis is bald* and (b) this theory can be grounded in an independently plausible theory concerning implicatures of compound sentences. Along the way we will say a lot about conditionals whose antecedents typically carry Gricean implicatures – these will be an important data source. Reflecting on these conditionals will help explain some odd data concerning vagueness, but it will also provide an interesting perspective on some problems concerning conditionals, such as Vann McGee's apparent counterexamples to *modus ponens*. We will also consider whether this explanation of (3) and (4) can be adopted by non-supervaluationists. The answer here will be that theorists who adopt non-classical logic almost certainly cannot adopt this solution, while theorists who retain classical logic but provide non-semantic theories of vagueness (such as, notably, epistemicists) probably cannot adopt the solution, though the evidence here is more equivocal. First, though, we shall survey the possible responses to the data about (1) through (4).

8.1. Famous Answers

Faced with such the challenges posed by these responses, theorists of vagueness seem to have five (or maybe six) options open to them.

Option One – Deny the Data

The simplest thing to do philosophically would be to deny the data; deny, that is, that there really are a substantial number of speakers who are willing to assent to (4), but not to (1), (2) or (3). Maybe after a substantial empirical investigation, this will turn out to be the right thing to say. But I doubt it is true. A poor reason for this is introspection. (Though as Jackson (1998: 37) notes, when philosophers say, 'Intuitively, p ', where p might be a proposition to the effect that such-and-such an example is or is not a case of knowledge, or causation, or justice, or whatever, the only evidence they usually have that p

really is intuitive is their own intuitions, and perhaps those of a few colleagues or students. And my intuitions about whether speakers in general are disposed to assent to certain sentences is a better guide to the facts than my intuitions about causation, knowledge or justice, because in the former case, but arguably not the latter, my intuitions are partially constitutive of the facts, since I am one of the speakers in question.)

A better reason is that it seems to be a fairly widespread assumption among experts in the field that the data is roughly as I have presented it. Some authors have explicitly asserted that this is the data (e.g. Burgess and Humberstone 1987 and Tappenden 1993), and others have implicitly conceded the same thing. Theorists who reject the law of non-contradiction typically feel they have some explaining to do. See, for instance, Machina (1976: 183-5), Tye (1994: 194) and Parsons (2000: 71) for acknowledgements of this and attempts at explanation. On the other hand, some of those who accept the law of excluded middle similarly feel an explanation is needed. See, for instance, Keefe (2000: 164), who proposes an similar, though less wide-ranging, explanation to the one I will provide below. The reason such theorists feel this way, I imagine, is that they note that we intuitively do assent to (4) but not (3), so they have to explain their divergence from ordinary practice. I will from now on assume that speakers do have these intuitions, though of course this is an empirical assumption, and much of the argument in what follows would lose some force if there was serious evidence against this assumption.

Option Two – Deny that the Data is Relevant

There is an obvious reason we might think that the data about assent to, and dissent from, various sentences is relevant to the theorist of vagueness. Such a theorist is in a position similar in broad respects to Quine's radical translator (Quine 1960: 26-35), though with two salient differences. First, she is trying to 'translate' her own language. This might not be a dramatic difference if one thinks that children learn their native language by a process similar to that which the radical translator learns the foreign language, though such an assumption seems to be rather implausible these days (Laurence and Margolis 2001). Secondly, she is *not* taking for granted that the logic and semantics of the language under investigation are classical. Still, the similarity is close enough that we should take native dispositions concerning assent to various sentences in various situations to be important data. But, it might be objected, we do not take untutored dispositions to be particularly important here. What really matters to our project are the *reflective* dispositions of speakers, and, it might be argued, speakers will not keep the dispositions described above at the end of the process of coming to reflective equilibrium. This is a more serious option than the first, and it requires a more subtle response. In a nutshell, the response I will give is that the theory I develop in this paper not only predicts but justifies speakers assenting to (4) but not (1), (2) or (3), and hence these dispositions can be kept in equilibrium. How good a response this is cannot be assessed without seeing my theory, so I will say no more about this until the theory is presented. (As a few authors have stressed, for example Sanford (1976), Tye (1990) and Tappenden (1993), there is also an argument that we should not assent to (3), based on the intuition that if a disjunction is true then there must be an answer to the question, "Well, which of its disjuncts is true then?" While this argument can be directly challenged, and has been by Dummett (1975), it nevertheless provides some reason for thinking that the intuition that (3) cannot be properly asserted will survive into equilibrium.)

Option Three – Radical Semantic Change

One could hold that the reason that the only one of the numbered sentences to which speakers assent is (4) is because (4) is the only one of them that is true. A sufficient motivation for holding such a view would be believing that (a) speakers are competent judges of the truth value of sentences such as (1)

through (4) and (b) they assent to such sentences iff they are true. As we discussed in chapter six, this option is taken completely by Burgess and Humberstone (1987), and is adopted in part by many other theorists. Defenders of ‘many-valued’ logics accept that (1), (2) and (3) are not completely true, though neither are they completely false. Supervaluationists accept that (1) and (2) are not true, and (4) is true, though (3) is also true, so there must be some alternative explanation for why speakers decline to assent to it. As several writers have pointed out, most notably Field (1986) and McGee (1991: ch. 4), the philosophical justification for such a move is somewhat dubious. In most fields of study, if there is some clash between our theories, the data, and classical logic, then what generally goes is our theory, unless we have good reason to impugn the data. It is very unlikely that the best move will be to dismiss classical logic, unless there are no other moves available. So the success of this option depends on the non-viability of other options, and I demonstrate below that a rival option is viable. (Also there are serious internal difficulties with taking the data at face value, a point we noted above that Burgess and Humberstone themselves do a pretty good job of presenting.)

Option Four – Moderate Semantic Change

Russell (1923) did not discuss (4), but agreed that (1), (2) and (3) might all fail to be true. This was not because he had a radically non-classical logic. Rather, it was because he thought that logic only applied to logically perfect languages, and natural languages are not logically perfect because they contain vague terms. While this position is not vulnerable to exactly the same methodological objection as option three, it does seem unhappy for two reasons. First, it is a seriously incomplete theory unless it tells us what kinds of reasoning we are allowed to use in natural language. Since instances of the law of excluded middle are not true, any argument that law as a conclusion must be flawed in some way, but Russell does not provide a systematic way to locate such errors, and no one developing his theory has done so either. Secondly, the theory must provide a way to provide truth conditions to sentences such that (3) is not true, and either (4) is true or we have an explanation of why speakers are disposed to assent to it even though it is not true. The first option here seems to lead to the difficulties that Burgess and Humberstone face, and the second is a theory schema in need of completion. In neither case does it seem this Russellian option is preferable to the position I will presently describe.

Option Five – Radical Pragmatics

We know for many sentences that whether speakers are disposed to assert them, or even assent to them, depends on many factors beyond the mere truth conditions for the sentence. In each of the following cases⁴, speakers may only assent to the sentence marked (a) if the condition marked (b) is not satisfied.

- (5) (a) That looks like a knife and fork.
(b) That is not a knife and fork.
- (6) (a) He drove carefully down the street.
(b) He used *reasonable* (as opposed to *unreasonable*) care in driving down the street.
- (7) (a) Her action was voluntary.
(b) Her action was blameworthy.
- (8) (a) Katie had several drinks and drove home.
(b) Katie had several *alcoholic* drinks and *shortly afterwards* drove home.

⁴ All borrowed, more or less literally, from Grice (1989).

In all cases, the condition mentioned in (b) is no part of the truth condition of the sentence (a)⁵, though it may be required for ordinary speakers to be willing to assent to that sentence. The pragmatic interpretation of (1) through (4) is that just as in (5) through (8), there is some difference between the situations⁶ in which speakers would willingly assent to the sentences, and the situations in which they are true. In what follows I will defend this option. We can, without assuming much at all about what the true theory of vagueness looks like, develop a pragmatic theory that predicts (and, for that matter, justifies) speakers assertoric practices concerning sentences like (1) through (4), and concerning a few more interesting cases that will be discussed below. While the existence of such a theory does not *entail* that various theories of vagueness based on non-classical logic are mistaken, indeed the pragmatic theory I sketch will, when combined with such theories generate true and interesting predictions, just as it would when combined with more conservative theories of vagueness, it does undercut the support for theories based on non-classical logics at a crucial point.

There is, perhaps, a sixth option available, which is to mix and match between the above accounts. Just how reputable this option is depends on just how systematic the mixing and matching is. One might claim that some of the dispositions under consideration will not be preserved in equilibrium, others can be explained pragmatically, and others are good guides to the semantics. If this is done unsystematically, then it is obviously philosophically dubious. Later in the paper I will suggest that some recent arguments against various theories of vagueness commit just this sin. But for now we will focus on the version of option 5 outlined in the introduction.

8.2. Truth, Assertion and Compound Sentences

Consider again sentence (8), which we will focus on for a while.

(8) Katie had several drinks and drove home.

The truth conditions for this sentence should be clear enough, though perhaps a little vague at the fringes.⁷ The sentence is true iff each conjunct is true. That is, (8) is true iff it is true that Katie had several drinks (in the time pragmatically specified as being under consideration) and drove home (again in *that* specified time). Assuming that the context specifies that the time under consideration is last night, then (8) is true iff Katie had several drinks last night and Katie drove home last night.

The sentence cannot be properly asserted, and speakers would not normally be disposed to either assert it or assent to it, unless Katie drove home shortly after having the said several drinks, and that the drinks in question were alcoholic. The reason it cannot be properly asserted unless this condition is satisfied is that hearers will normally conclude from the existence of the utterance that the conditions are satisfied, and hence the speaker would mislead the hearers if they were not. There is one

⁵ Contra the suggestions of Wittgenstein (1953), Hart (1961) and Ryle (1949) in the cases of (5), (6) and (7).

⁶ In using this term I do not mean to endorse *all* of the details of the views of Barwise and Perry (1983); I merely use it as the least loaded term available in the circumstances. If one so desires, one can understand 'situations' to be centred possible worlds in everything that follows.

⁷ The vagueness will not be directly relevant *here*. For most (but not all) of the points I want to make below, we could instead use (8')

(8') Katie drank a bottle of scotch and drove home.

It is convenient to have a sentence that does not *say* that the drinks were alcoholic, so we will stay with (8) for now. I am grateful to a conversation Peter Smith for clearing up some of the possible confusions here.

other condition that must be satisfied before speakers will happily assert (8). They must know (or at least take themselves to know) that all the conditions mentioned above are true. So for (8) to be assertable, a certain fact about the world must be true, Katie must have had several alcoholic drinks and shortly afterwards drove home, and a certain fact about the speaker must be true, she must have a justified belief that the fact about the world is true. In general (though perhaps not always) we will be able to make such a division into facts about the world that can be reasonably assumed to be communicated by an utterance and hence must be true before a sentence can be properly asserted, and facts about the speaker (usually that they have a justified belief that those facts about the world hold) that must also be true before that speaker can properly assert the sentence. No doubt there will be practical difficulties in any case in making this division, and in some cases there may even be conceptual difficulties in carrying out this task. (We will come back to this point below.) Recognising this difficulty, we will for now carry on as if the division can be made. For a sentence S , say the *semantic content* of S is the set of situations in which S is true, the *objective pragmatic content* of S is the set of situations such that the conditions about the world necessary for S to be asserted are satisfied and the *subjective pragmatic content* of S for x is the set of situations in which S justifiably believes that objective pragmatic content of S is satisfied. The idea is that x should be happy, on reflection, to assent to S in just those situations in the subjective pragmatic content of S . We will write $T(S)$ for the semantic content of S , $O(S)$ for its objective pragmatic content, and $A(S, x)$ for its subjective pragmatic content for x . The semantic content of the sentence on an occasion is what Grice said was *said* (in his favoured sense) by uttering the sentence, while the objective pragmatic content is what he said is *implicated* (Grice 1989: 118). Stanley and Szabo (2000: 230) use ‘communicated’ here, which is probably more perspicuous in virtue of being less technical. (Note that when I say that the semantic content is an unstructured entity, a set of situations, I do not rule out the possibility that the sentence has that content in virtue of expressing a structured proposition.) The subjective pragmatic content corresponds rather closely to Quine’s affirmative stimulus meaning (Quine 1960: 33).

It is commonly assumed that semantic content must be compositional. (Though this assumption is not universal; see Schiffer 1987 and McGee 1990.) This assumption may or may not be true, but there is some evidence that objective pragmatic content is compositional. (Indeed, this is an important reason for recognising objective pragmatic content as well as subjective pragmatic content.) Consider the indicative conditional (9).

(9) If Katie had several drinks and drove home, then she broke the law.

It seems that $O((9))$ includes all the situations we might find ourselves in these days. Given that there are laws against driving while intoxicated, and that the antecedent implies that Katie drove intoxicated, we are happy to assent to that conditional. (Perhaps things would be different if we somehow knew that Katie was immune to motor laws, but let us set that aside.) But even though intuitively $O((9))$ includes all situations we might hope to find ourselves in, there is a good *argument* that $T((9))$ does not include some salient situations in everyday life. Consider the world in which last night Katie drove home sober, then had several drinks, and broke no laws for the evening. Then (9) is a conditional with a *true* antecedent and a *false* consequent. And that indicative conditionals with true antecedents and false consequents are false is the closest thing there is to a point of consensus in theories about conditionals. Of course not quite everyone joins the consensus, most prominently McGee (1985). Still, the principle is called the Uncontested Principle by Jackson (1987), so my claim that it is a consensus is not exactly idiosyncratic. So $T((9))$ does not include some situations that *are* included in $O((9))$. Further, we can work out what $O((9))$ is without knowing what $T((9))$ is; even if doubt about various semantic theories

concerning indicative conditionals means that we are unsure of the truth conditions for (9), we know that $O((9))$ includes all the situations we are likely to encounter. This explains why we can assert (9) while knowing next to nothing about Katie; given reasonable background assumptions, its objective pragmatic content is more or less trivial.

These considerations decisively refute one possible theory of how we calculate objective pragmatic content, a theory that Grice seems to take to be true.⁸ On this hypothesis, we calculate the objective pragmatic of a sentence by first consulting our linguistic knowledge to determine its semantic content, then employing our mastery of the Gricean maxims to work out what conversational implicatures it might have, and the objective pragmatic content is the set of situations in the semantic content where those implicatures are all true. This can't be right in detail, for it predicts $O((9))$ will be a subset of $T((9))$, which we have seen is not true. And it can't even be right in broad outline, because it predicts we should be unsure of the objective pragmatic content of a sentence until we know its semantic content. This seems clearly untrue in the case of (9); we can know its objective pragmatic content while remaining quite unsure of its semantic content. (Stanley and Szabo (2000: 231) say that "Cases where the speaker knows the proposition communicated without the proposition expressed ... are highly exceptional." The above considerations seem to suggest either that indicative conditionals, or at least indicative conditionals whose antecedents typically carry Gricean implicatures, are exceptional cases or that Stanley and Szabo's claim is mistaken.)

Those considerations suggest that the objective pragmatic content of compound sentences is tied less closely to the semantic content of that sentence, and more closely to the objective pragmatic content of its constituent sentences. In light of this suggestion, the following hypothesis, called the Compositionality of Objective Pragmatic content, or COP, thesis, might be plausible.

COP The objective pragmatic content of a compound sentence is a function of the objective pragmatic contents of its constituents, with the function given by the operator or connective used to form the compound.

When stated in full generality like that COP is a bit obscure, but it becomes clear with a few examples. COP entails that $O(\text{If } A \text{ then } B)$ is $\text{If } O(A) \text{ then } O(B)$. In the case of (9), this is exactly the right answer. Further, it entails that $O(A \text{ or } B)$ will be $O(A) \text{ or } O(B)$ and $O(\text{Not } A)$ will be $\text{Not } O(A)$. Again, consideration of sentences where (8) is embedded in various sentences suggest that COP is on the right track.

One might worry that this one example can hardly support a theory as wide ranging as COP. This of course is true; a large part of the argument for COP is that it generates a theory that makes such surprising true predictions when applied to vagueness. (See, in particular, the discussion of complex contradictions in section 4 below.) But it is worth noting that the points about (8) and (9) above do not just turn on features to do with the ordering implication in conjunctions. Kent Bach (1994: 134) notes that one can often use (10) to sooth someone who is in a relatively mild state of disrepair. (Imagine a mother saying this to an injured child, or a doctor reporting good test results to a patient.)

(10) You're not going to die.

Normally, when this is uttered, the speaker knows, or should know, that it is literally false. Of course the person addressed is going to die some time, so the semantic content of (10) is false. This reasoning

⁸ That last claim is contentious. Grice only needs the premise that things are *as if* the theory I am about to describe is true, and even if that theory is false, it is possible that the relevant things are as if the theory is true.

is not conclusive; the sentence might be elliptical and it might be clearly true once the ellipsis is completed. But the more plausible position seems to be that the sentence is false. The speaker communicates that the hearer is not going to die soon, or from the particular illness they are suffering, and that is the objective pragmatic content of the sentence. And we can note that this content seems to be carried over into conditionals. Jack is working on a major project, and his manager Jill is concerned he is taking too much time off with illness. While he is at home with one minor illness, Jill emails him the following directive.

(11) If you're not going to die, then you should be in at work.

Heartless, perhaps, but the intended message is clear. If Jack is not going to die from this particular illness, or at any rate in the near future, he should be at work. Jack could hardly say that this conditional directive (threat?) did not apply to him, because as a mortal he is sure to die. COP is not entailed by two examples any more than it is by one, but it is worthwhile noting that we have not had to rely on particular features of (8) or (10) to support COP.

If something like COP is correct, then it is important to distinguish between objective and subjective pragmatic content.⁹ Note that if we replaced *objective* with *subjective* pragmatic content in COP, generating a thesis we may call CSP, we get a clearly false thesis. It is not the case that we are only happy to assert *A or B* when we are happy to assert *A* or we are happy to assert *B*. We may assert *Either X will be the next Prime Minister or Y will be the next Prime Minister*, for suitable *X* and *Y*, when we don't know who the next Prime Minister will be, but are very confident that it will be *X* or *Y*. Indeed, we might *only* assert it if we don't know who the next Prime Minister will be; if we did know this we would assert it rather than just the disjunction. So CSP is false, but this does not show that COP is false.

8.3. Application to Vagueness

On most theories of vagueness, if *F* is a vague predicate, then we can distinguish between *a* being *F*, and *a* being determinately *F*. And, again on most theories, the conditions under which it is proper to say that *a* is *F*, or to assent to the claim that *a* is *F*, are those where *a* is determinately *F*. On the epistemic theory of vagueness, we can only say that *a* is *F* if it is known that *a* is *F*, and, according to that theory, that means that *a* is determinately *F*. (What makes the epistemic theory of vagueness epistemic is that it interprets 'determinately' as an epistemic operator.) On the supervaluational theory, we can only properly say that *a* is *F* if *a* is *F* on all¹⁰ precisifications, and that is what it is for *a* to be determinately *F* on that theory. On degree of truth theories, we can assert that *a* is *F* iff *a* is *F* to a very high degree, and that is what it is for *a* to be determinately *F* on that theory. In short, the pragmatic content of $\lceil a \text{ is } F \rceil$ is that *a* is determinately *F*, or, for short, $\Box Fa$. Note that on the epistemic theory, this is the *subjective* pragmatic content, while on the supervaluational and degree of truth theories it is the *objective* pragmatic content. This reflects the differences between the ways the theories understand determinate truth. This point acquires some importance soon, so we will return below to what epistemicists might take the objective pragmatic content of $\lceil a \text{ is } F \rceil$ to be. For now we will focus on those theories where $\Box Fa$ is the

⁹ I am indebted here to conversations with Tim Maudlin and Brian McLaughlin.

¹⁰ Or perhaps most. Little will turn on this here, but the ability of supervaluational theories to handle various paradoxes (such as the Sorites and the problem of the many) might depend on just how this principle is framed. For salient discussion on this point see Lewis (1993).

objective pragmatic content of $\lceil a \text{ is } F \rceil$. On those theories, COP predicts that objective pragmatic content of (3) and (4) will be (12) and (13)

- (3) Louis is bald or Louis is not bald.
 (12) $\Box(\text{Louis is bald}) \text{ or } \Box(\text{Louis is not bald})$.
 (4) It is not the case that Louis is bald and that he is not bald.
 (13) It is not the case that $\Box(\text{Louis is bald})$ and that $\Box(\text{Louis is not bald})$.

Note that on almost any theory of vagueness one cares to consider, if Louis is a penumbral case of baldness, then (12) will be false and (13) true. (12) is false because Louis's penumbral status makes both $\Box(\text{Louis is bald})$ and $\Box(\text{Louis is not bald})$ come out false. Hence if COP is true, speakers should decline to assent to (3), but should assent to (4). Since they do, this is good news for COP.

We have not said how COP should apply to quantified sentences. These are a little harder to incorporate into the theory than sentences that are formed by familiar propositional connectives because the logical form of these sentences is less transparent. I will assume¹¹ that quantified noun phrases are restricted quantifiers. Hence the logical form of (14) will be (15).

- (14) $Q \text{ Fs are Gs}$
 (15) $[Qx: Fx] Gx$

Although COP as it stands is silent on quantified sentences, we can naturally generalise it. And the natural thing to say, given COP, is that if the logical form of (14) is (15) then its objective pragmatic content is (16).

- (16) $(Qx) [\Box Fx : \Box Gx]$

This lets us explain one of the consequences of supervaluationism that is, if anything, more surprising than its endorsement of the law of excluded middle. The following presentation of the puzzle is due to Jamie Tappenden. Let $P(n)$ abbreviate 'A man with exactly n cents is poor'. "Since the supervaluation deems the conditional premise of the sorites paradox false, it deems true the claim that there is an n such that $(P(n) \ \& \ \neg P(n + 1))$. Alas there is no such number." (Tappenden 1993: 564) Let us abbreviate further, and say that n is the poor borderline iff it is such that $(P(n) \ \& \ \neg P(n + 1))$. Then the dubious claim is (17), written symbolically as (18), and Tappenden's alternative judgement is (19), or in symbols (20).

- (17) Some number is the poor borderline.
 (18) $[\exists x N(x)] (PB(x))$
 (19) No number is the poor borderline.
 (20) $\neg([\exists x N(x)] (PB(x)))$

From what we said above, it follows that the objective pragmatic contents of (17) and (19) are (18a) and (20a).¹²

¹¹ Following Barwise and Cooper (1981), Higginbotham and May (1981) and, most directly, Neale (1990)

¹² Nothing turns on it here, but whether the box at the front of (20a) should be there depends on just what the logical form of *No Fs are Gs* is. I have assumed, without any good reason, that it is $\neg[\exists x: Fx](Gx)$, but other interpretations are possible. Nothing turns on it because however we write (20), (20a) will be true.

- (18a) $[\exists x \Box N(x)] (\Box PB(x))$
 (20a) $\Box \neg([\exists x \Box N(x)] (\Box PB(x)))$

Since (18a) is false, and (20a) true, we have an explanation not only of why one cannot properly assert (17), but of why one can assert its negation, (19).

So our hypothesis, derived from COP, can explain a few puzzling pieces of data. And if COP is generally correct, then it does so in a fairly systematic way. Tappenden suggests that we endorse (4) but not (3) because “Noncontradiction in these cases is a “no overlap” condition” while “excluded middle functions as a “sharp boundaries” condition.” (565) No doubt this is true, but it would be nice to have a systematic explanation of why this is so, and COP promises to give us one.

We notes above that the *subjective* pragmatic content of compound sentences is not built from the subjective pragmatic content of its components in the way that objective pragmatic content is built up. This means that if the *definitely* operator in the pragmatic content is part of the subjective pragmatic content, we cannot give the above explanation for the unattractiveness of (3). This is not a problem for a supervaluational theory, but it is a problem for the epistemic theory. When Louis is a penumbral case of baldness, the reason we are happy to neither assert that he is bald nor assert that he is not is that we do not know which. There is a close analogy, according to this theory, between our unwillingness to make these assertions and our unwillingness to assert either that Louis was born in France or that he was not when we do not know which. (I assume there is no vagueness concerning where Louis was born, though of course there is a faint possibility that this would be vague. The analogy is both drawn and alleged to be supported by data in Bonini *et al* 1999.) But it seems there is an important disanalogy between these two cases; namely, in the case of ignorance we are happy to say that the relevant instance of excluded middle is true, whereas our assessment of (3) is, as Tappenden puts it, “range from mixed to strongly negative” (1993: 565).

There is a way out of this difficulty for the epistemicist, though it does require relaxing the analogy between cases of vagueness and traditional cases of ignorance. We said above that the subjective pragmatic content of a sentence (on an occasion) was that the speaker knew the objective pragmatic content was satisfied. In effect, we took objective pragmatic content as primary, and subjective pragmatic content is derived from it. We could have done things the other way around. Loosely following Quine’s lead, we will take subjective pragmatic content as primary, and say that the objective pragmatic content is what is common to subjective pragmatic content all (or perhaps most) occasions the sentence is used. If part of the subjective pragmatic content on every occasion of utterance of a simple sentence is that the speaker knows the sentence is true, then it will follow that part of the *objective* pragmatic content is that the sentence is *knowable*. So we can interpret \Box in each of the statements of objective pragmatic content as a *knowability* operator, and we get all the right results: (10) and (18a) turn out to be false, (11) and (20a) turn out to be true.

If true this account would explain the data, but its plausibility seems open to question. The theory says that objective pragmatic content of a simple sentence is what is common to all the different subjective pragmatic contents, and then says that the objective pragmatic contents of complex sentences are composed out of the objective pragmatic contents of the components, and says this fact explains our reactions to (3) and (4). But that fact (if it is a fact) stands in need of explanation just as much as our reactions to (3) and (4) do. If objective pragmatic content is something as artificial as this, it seems mysterious why it should so neatly compose. If, on the other hand, objective pragmatic content is something that speakers understand in virtue of understanding the sentence (and perhaps even something they understand before understanding the semantic content), then we have an explanation

for why it is compositional: it has to be if speakers' understanding of the language is to be productive. So while epistemicists *can* adopt something like COP as an explanation for our reactions to (3) and (4), their adopting of it must rest on premises that seem surprising, and possibly in need of explanation.

8.4. Modifying COP

If our only data was that people are hesitant about instances of the law of excluded middle, like (3), but are not resistant to asserting simple instances of the law of non-contradiction, like (4), then COP would do an excellent job in explaining the data that we have. But this does not seem to be all the data we have. In two ways our willingness to assert sentences goes beyond what is suggested above. In this section I will set out that data, and then suggest a natural revision of COP that explains this data, as well as the data presented above.

Above we have stressed what Jamie Tappenden calls the truth-functional intuition. This intuition, directly or indirectly, causes us to resist disjunctions when we know that for each disjunct there is no fact of the matter as to whether it is true. As Tappenden notes, though, our intuitions towards vague sentences are also guided at times by a penumbral intuition. When under the sway of this intuition, we tend to judge disjunctions as true if all the possibilities other than the disjuncts in question have been ruled out. To take a classic example from Fine (1975), if a shade of colour is around the border between red and orange, then in the right frame of mind we might be prepared to say *That is red or orange*. To be sure, as Tye (1990) and Tappenden (1993) note, this intuition can waver in the face of sustained argument, such as the forceful suggestion that if a disjunction is true there should be a fact about which disjunct is true. And as Machina (1976) makes clear, some philosophers do not feel this intuition at all.

I think the penumbral intuition is a real, widespread phenomena, and it is incumbent on a theory of vagueness to explain it. My main reason for regarding it this way, despite the comments of Tye, Tappenden and Machina, is the prevalence of the penumbral intuition in some of the social sciences. This was argued for above (see page 69), so we will not repeat it here, save to note the conclusion that social scientists are prepared to *talk* in a certain way, accepting disjunctions even when they know they could not, in principle, have reason to accept either disjunct, and that this is an important piece of data that theorists of vagueness must explain. The data is not of a different kind to what had been previously considered, but it does reinforce the claim that the penumbral intuition must be accommodated.

One other point from chapter 6 is important here. We noted there that speakers' reluctance to endorse simple contradictions did not extend to complex contradictions. So while (21) and (22) are definitely odd, both (23) and (23a) can be uttered if Louis, the king of France, is known to be a borderline case of baldness.

- (21) ?Louis is bald and Louis is not bald.
- (22) ??Louis is bald, and the King of France is not, and Louis is the King of France.
- (23) It is not the case that Louis is bald, but nor is it the case that he is not bald.
- (23a) It is not the case that Louis is bald, but nor is it the case that His Majesty is not bald.

Some contradictions, such as (23) can be properly asserted on account of vagueness, and this is not due to their being idiomatic. In the terminology of the previous section, the objective pragmatic content of (23) is not the null set, it is the set of situations where Louis is a penumbral case of baldness. This striking piece of data needs to be explained. And it should be immediately clear that the explanation cannot be that (23) is true. After all, (23) is a contradiction, and contradictions have a tendency to be false. So the explanation for it must be pragmatic. As it stands, COP cannot provide that explanation.

But a small alteration to COP can do so, and when combined with the right kind of theory of vagueness, can explain why we are happy to accept disjunctions without a determinately true disjunct, at least while in the social science classroom. I will first state the theory, called POP (loosely for Pragmatic determination of Objective Pragmatic Content), then explain how it applies. I will first state the special case of POP for sentences that have no differences between their semantic and objective pragmatic contents other than those caused by vagueness. This special theory will be called POP_v.

POP_v Let S be a sentence that has no differences between its semantic and objective pragmatic contents other than those caused by vagueness. Then there is a sentence S' generated by adding \Box operators to S so that every term in it apart from sentential connectives is inside the scope of a \Box operator, and $O(S) = T(S')$. Which such sentence S' satisfies this condition on an occasion where S is used is determined by pragmatic features of utterance and occasion.

So, for example, if S is $\lceil a \text{ is } F \text{ or } b \text{ is } G \rceil$, then POP_v says that the objective pragmatic content of S is either $\Box Fa \vee \Box Gb$ or $\Box(Fa \vee Gb)$. More generally, any sentence you can generate from S by adding boxes so that every part of S , except the sentential connectives, is inside the scope of a box, could be the objective pragmatic content of S . Letting l be a name for Louis, and B refer to the property of baldness, then POP_v says that the objective pragmatic content of (3) is (24a) or (24b), and that the objective pragmatic content of (4) is one of (25a) through (25c).

(3) Louis is bald or Louis is not bald.

(24a) $\Box Bl \vee \Box \neg Bl$

(24b) $\Box(Bl \vee \neg Bl)$

(4) It is not the case that Louis is bald and Louis is not bald.

(25a) $\neg(\Box Bl \wedge \Box \neg Bl)$

(25b) $\neg\Box(Bl \wedge \neg Bl)$

(25c) $\Box\neg(Bl \wedge \neg Bl)$

In each case, it says there should be some indeterminacy in precisely how the objective pragmatic content is generated. But, as we said in the introduction, a crucial difference arises between the two cases if we assume a broadly supervaluational (or epistemic) interpretation of \Box . In (3), POP_v predicts that depending on the context, the objective pragmatic content will either be (24a), which is false, or (24b), which is true. So it predicts that whether (3) is assertable will depend on the broader features of the context that select which of these will be the objective pragmatic content. In (4), POP_v predicts that whatever method is pragmatically selected to generate the objective pragmatic content, it will be true. So (4) should be always assertable. In each case, we seem to get a rather pleasing correlation between predictions and data.

The assumption here that \Box would get a supervaluational (or epistemic) interpretation is crucial. Interpret $\Box A$ as meaning A has a high truth value, in a typical degree of truth theory, and we do not get the conclusion that (3) should sound trivial in some contexts, since (24b) is not guaranteed to be true, and nor do we get the conclusion that (4) should always sound trivial, since (25c) is no longer guaranteed to be true. In fact, if Louis is anything like a penumbral case of baldness, (25c) is guaranteed to be false on those theories. So if proponents of that theory want to explain why (4) sounds trivial, they must appeal to something other than POP_v. Supervaluationists, however, can explain the data just via POP_v. So can epistemicists, provided they can discharge the burden outlined earlier of explaining

how a subjective feature like knowability can make its way into objective pragmatic content. From now on we will assume, with the supervaluationists and epistemicists, that all classical tautologies are true, and all classical anti-tautologies are false.

Most surprisingly, POP_V is consistent with the hypothesis that complex contradictions, like (23) should be properly assertable, and hence have a non-degenerate objective pragmatic content. POP_V predicts that the objective pragmatic content of (23) is one of (26a) through (26d).

(23) It is not the case that Louis is bald, but nor is it the case that he is not bald.

(26a) $\neg \Box B \wedge \neg \Box \neg B$

(26b) $\Box \neg B \wedge \neg \Box \neg B$

(26c) $\neg \Box B \wedge \Box \neg \neg B$

(26d) $\Box(\neg B \wedge \neg \neg B)$

(26b) through (26d) are all false, since they are all inconsistent, or entail inconsistencies by obvious steps. But (26a) is not inconsistent, indeed it is true. So POP_V explains why a contradiction, like (23) can be used to convey a true message, despite not being idiomatic. This is a rather surprising piece of data to have, and it is even more surprising to have a neat explanation of it. For what it is worth, and depending on your preferred philosophy of science it might be worth a lot, when I was developing this theory, the explanation of (23) appeared as a prediction, not a retrodiction. I had no idea that there could be non-idiomatic contradictions which, because of vagueness, could be used to convey true messages, until I realised POP_V predicted the existence of such sentences, and then I realised (23) was such a sentence. If one tends to value surprising and true *predictions* higher than any retrodictions, this little bit of autobiography has epistemic importance. (It does seem rather surprising to think that autobiographical facts like this could effect the plausibility of a theory; so surprising in fact that one might view it as a problem for the prediction/retrodiction distinction. But that is a matter for a book far removed from this one.)

So POP_V can explain some data that COP cannot. Still, we have insufficient reason to switch to POP_V unless it can be embedded in a theory of the same level of generality as COP. To that end, I set out the general version of POP. For any compound sentence, there are going to be several ways to partition the sentence into sub-sentences that one can 'treat as simple'. To treat a sub-sentence as simple, in the sense intended here, is to not take into account the fact that the sentence has sentences as parts when evaluating its objective pragmatic content. For a sentence that one treats as simple, the objective pragmatic content is generated from the form and semantic content by the application of broadly Gricean maxims, including, if this is not included already, a maxim that one should not assert sentences unless they are *determinately* true. Sentences that have no sentences as proper parts in their logical form have to be treated as simple.¹³ For other sentences, though, there are often choices as to which sentences can be treated as simple for purposes of generating objective pragmatic content. For example, if S is $\lceil \text{If } S_1 \text{ and } S_2 \text{ then } S_3 \rceil$, then one can treat S_1 , S_2 and S_3 as simple sentences, or $\lceil S_1 \text{ and } S_2 \rceil$ and S_3 , or, possibly, S itself. (The 'possibly' here is because it is unclear whether we can treat a conditional as simple, just because we can only treat as simple sentences that we (at least tacitly) know

¹³ When we are just investigating the objective pragmatic content of conditionals and disjunctions this condition is relatively easy to interpret. Some complications arise because POP, like COP, is intended to apply to quantified sentences as well, and indeed mimic COP's explanation of why (17) is not assertable though (19) is. The intent here is that the Gx in $\lceil \exists x: Fx \rceil Gx$ should count as a sentence for purposes as POP.

the truth conditions of, and we are so in the dark about how conditionals work that it is not clear this is possible.) Then POP says:

POP The objective pragmatic content of a compound sentence is a function of the objective pragmatic contents of its sub-sentences that are treated as simple, with the function given by the operator or connective used to form the compound. The choice of which sub-sentences are treated as simple is determined by the syntactic features of the sentence and the context. The objective pragmatic content of sentences treated as simple is determined by a direct application of broadly Gricean rules.

The reference here to syntactic features is because there seem to be some quite general relations between surface-level syntax and which sentences are treated as simple in the default evaluation of a sentence. In general, the more variables that are relevant to the objective pragmatic content of two sentences and not determined by the surface syntax, the less likely those sentences are to be treated as simple. Conversely, the fewer 'hidden variables' shared by two sentences, the more likely they are to be treated as simple. We can see this illustrated by an example we have already examined. For ease of exposition, we will use a slight variant on the example discussed above.

(27) If Katie had several drinks and she drove home, then she broke the law.

In a normal utterance of (9), we evaluate each clause as being about a specific time. For example, even if there was a time ten years ago when Katie had several alcoholic drinks and drove home shortly afterwards, we would not normally consider the antecedent of the conditional satisfied unless she repeated this behaviour more recently. (There is an obvious exception to this if the events of ten years ago are for some reason under consideration.) The sense of 'satisfaction' here is meant to be rather pragmatic; it is the sense in which the antecedent of (9) is not satisfied if Katie drank water all night after driving home sober. So in the pragmatic content of the antecedent of (9) there is a suppressed reference to a time period. We can make this explicit, as in (28), without noticeably changing the pragmatic content.

(28) If, last night, Katie had several drinks and she drove home, then she broke the law.

In (28) we added the suppressed time reference to the conjunction that is its antecedent. This addition did not, it seems, change the pragmatic content. If, however, we add the suppressed time reference to both conjuncts, we do get a change in pragmatic content.

(29) ?If Katie had several drinks last night and she drove home last night, then she broke the law

To my ear, at least, (29) does not seem as trivial as (28), and certainly not as trivial as (27) or (9). POP has an explanation for this. Once the two conjuncts in effect stop sharing an unuttered constituent, there is less temptation to treat the conjunction as a simple sentence. If we treat the two conjuncts as simple, then the objective pragmatic content of the conjunction is just the conjunction of the objective pragmatic content of the conjuncts. And that means that the objective pragmatic content does not imply that the driving happened after the drinking, since this is not implied by the objective pragmatic content of the conjuncts, taken separately or together.

This hypothesis about which sentences are treated as simple has implications for vagueness. It predicts that the objective pragmatic content of (23) will be (23a), and hence that (23) is not just possibly assertable, but actually assertable. Returning to Fine's example of the colour patch that is somewhere between being perfectly red and perfectly orange, it predicts that (30) should sound better than (31).

- (30) The patch is red or orange.
 (31) The patch is red or the patch is orange.

I think this prediction is correct, though my intuitions here are perhaps getting unreliable.

The more important point is that there seems to be some evidence for POP outside of vagueness. Since POP also does an excellent job at explaining several puzzling features concerning vagueness, this suggests it is a rather well-supported theory.

To conclude this chapter, I will mention four points that follow from the discussion so far. First, I will note that everyone needs a theory something like POP if they are to explain the data, so the fact that supervaluationists (and their fellow travellers) need POP to explain some data is no cost to them. Secondly, I will note that the plausibility of POP reduces the plausibility of some arguments against various theories of vagueness. Thirdly, I will outline what POP tells us about Sorites arguments. And fourthly, I will return to the point of whether epistemicists can appeal to POP to explain the data under consideration here. I will suggest that when we look closely at Vann McGee's 'counterexample to modus ponens' we find some evidence to suggest that they can. The evidence is hardly conclusive, and POP still sits more comfortably with semantic than with epistemic theories of vagueness, but it may provide some solace to epistemicists.

8.5. The Need to POP

POP explains all the data we have, and makes surprising true predictions. But there might be other theories more deserving of our assent. In particular, it might be possible to provide a semantic explanation of the data we have seen so far, and if so, it *might* be preferable to accept such an explanation. It is not entirely clear whether such a semantic explanation of the data *would* be preferable if it were possible, because it has become fairly standard practice to prefer pragmatic explanations of data to semantic explanations. Grice, for instance, never conclusively proved that the semantic theories of Austin, Strawson and Wittgenstein that he attacked in the first of the William James lectures were false, just that they were unnecessary given his pragmatic theories. And most of us took that to be sufficient to make the case, presumably because there is a preference for pragmatic explanations in the kind of cases Grice considered. I will not stress this point here, because it seems all the semantic rivals to POP are provably inferior. I will only consider two rivals, because it shall quickly become clear that all such rivals face a fairly pressing problem.

The first rival theory to consider says that vagueness shows that the correct logic for natural language is intuitionist. Given that the original task we set ourselves was to explain why an instance of the law of excluded middle, (3), was not acceptable, while the matching instance of the law of non-contradiction, (4), was acceptable, intuitionism may seem to be a natural refuge. After all, intuitionism famously rejects excluded middle while accepting non-contradiction. It might be objected immediately that vagueness gives us no reason to drop the principle of double negation elimination. But it is not clear that this is so. After all, it is possible to read (23) as a denial of the conditional *If it is not the*

case that *Louis is not bald, then he is bald*, and as we argued above, it is possible to read (23) in such a way that it is acceptable.

The main problem facing intuitionists is that there seem to be acceptable vague sentences that are intuitionistically unacceptable. One simple example is (23). Even though intuitionists reject double negation elimination, and excluded middle, it is inconsistent to *deny* instances of either principle. But, as noted, it seems that (23) might be a denial of one or other of these principles. More tellingly, if we are to let surface structure be our guide as to semantic content, then we will say that (23) is a contradiction. And since all contradictions are false in intuitionist logic, intuitionists will counsel its rejection. Of course, we need not (and perhaps should not) say that surface structure shall be our infallible guide. So we might say that (23) is intuitionistically acceptable because it expresses something other than a contradiction, or that it is acceptable despite being a contradiction because there are pragmatic rules stating when false sentences are acceptable. Some such approach is clearly possible: if the intuitionist just buys POP then she can explain the data just as well as anyone else. But now intuitionism has ceased to be a rival to POP. Intuitionism plus POP makes for an intriguing theory of vagueness, but not one we need consider if we want to know whether POP is in the *best* explanation of the data.

The other semantic theory to consider here is the one set out by Burgess and Humberstone (1987). We noted above a few problems with that theory. Some of these, such as the failure of congruentiality, are not particularly relevant here. What are relevant are the difficulties their theory has with complex contradictions, such as (23). That sentence has the form $\neg A \ \& \ \neg\neg A$, and sentences of that form are always false on Burgess and Humberstone's theory. Moreover, just the kind of intuitions that would lead us to accept (4) but reject (3), that is, just the intuitions that Burgess and Humberstone's theory rests upon, lead to acceptance of (23). So it looks like Burgess and Humberstone's theory has not provided a systematic theory of all the intuitions it attempts to systematise. It also seems unlikely that their theory can explain why workers in social sciences feel free to ignore complications due to vagueness in their theories, whereas the theory based on POP can explain this.

More generally, it is hard to see how any semantic explanation of the data can account for these two phenomena. Since (23) is a contradiction, it is *false*, so we cannot possibly explain all the data here semantically. And the appropriate logic for natural language does not change when we start doing social science, but a semantic explanation of the phenomena can't explain how the phenomena are so contextually variable. Whether POP is precisely right or not, these considerations suggest strongly that it is in the right ballpark, and these semantic explanations are not.

8.6. POP and other arguments

The attraction of using many-valued logics in a theory of vagueness is that they provide an easy explanation of the unacceptability of (3). The big problem with such theories is that they provide no obvious explanation of the acceptability of (4). This, I think, is the motivation underlying Williamson's argument in the following passage.

The sentences *He is awake* and *He is asleep* are vague. According to the degree theorist, as the former falls in degree of truth, the latter rises. At some point they have the same degree of truth, an intermediate one ... the conjunction *He is awake and he is asleep* also has that intermediate degree of truth. But how can that be? Waking and sleep by definition exclude each other. *He is awake and he is asleep* has no chance at all of being true... Since the conjunction in question is clearly incorrect it should not have an

intermediate degree of truth...How can an explicit contradiction be true to any degree other than 0? (1994: 136)

It is, I suppose, noteworthy that Williamson did not make the following argument:

At some point the disjunction *He is awake or he is asleep* also has that intermediate degree of truth. But how can that be? Waking and sleep by definition exhaust the possibilities. *He is awake or he is asleep* has no chance at all of being false...Since the disjunction in question is clearly correct it should not have an intermediate degree of truth...How can an explicit instance of the law of excluded middle be true to any degree other than 1?

This alternative argument has, it seems, no persuasive force whatsoever. So whatever difficulty is being brought out by Williamson's argument must turn on the differences between his argument and the alternative argument. Hence the argument cannot just be, for example, that only sentences that are (completely) true in some possible situations and (completely) false in others can have intermediate truth values. If that principle were right, our alternative argument here would also go through. Nor can Williamson's argument just be that we intuit that contradictions have degree of truth zero, so contradictions have degree of truth zero. As epistemicists must accept, some sentences that intuitively have degree of truth less than one really do have degree of truth one. (Instances of the law of excluded middle seem to be good candidates.) For a similar reason, the argument cannot just be that we intuit that instances of the law of non-contradiction are true. Once we have accepted that theories can force us to revise intuitions about which sentences are true, we are in no position to insist that a particular theory must respect a particular intuition about the truth of various sentences.

The best form of Williamson's argument, I think, appeals to assertability. The argument cannot just be that *He is awake and he is asleep* are not assertable. Degree theorists agree that sentence is not assertable. Remember that their criteria of assertability is that a sentence is assertable iff it has a *high* degree of truth. Since *He is awake and he is asleep* has at most degree of truth 0.5, and 0.5 is not high, that sentence is definitely not assertable. If, however, we look at the negation of that sentence we do get an interesting disagreement. Intuitively, *It is not the case that he is awake and he is asleep* can be asserted, even though its degree of truth is merely 0.5. This, I think, is the strongest interpretation of Williamson's argument. What is wrong with *He is awake and he is asleep*, the reason it seems wrong to give it even a moderate degree of truth, is that its negation can be confidently asserted. If the degree-theorist cannot explain *this*, then her theory is in tatters. It is no good to say here that intuitions about assertability are not philosophically important; the whole motivation for the many-valued account is that it captures certain important intuitions about simple cases, in particular about instances of the law of excluded middle. So this is a serious challenge. But, as we have seen, it is a challenge that the degree-theorist can meet if she accepts POP. With POP, we can explain how it is that *It is not the case that he is awake and he is asleep* can be assertable even though it does not have a high degree of truth, because its pragmatic content is *true*.

Of course, similar arguments against epistemicism and supervaluationism, arguments that turn on the fact that sentences like (3) cannot be asserted, will also fail if epistemicists and supervaluationists can explain the unassertability of (3) using POP. Much of the motivation for degree theories, tracing back to early work by Goguen (1969) and Zadeh (1975), centred around the fact that we will not assert any of the sentences: *Louis is bald*; *Louis is not bald*; *Louis is bald or not bald*. It might be argued that if (3) were true then competent speakers would not be hesitant to assert it, and certainly would not deny it.

Hence theories that suggest it is true, such as epistemicism and supervaluationism, are mistaken. This reasoning is not without merit (it seems, for example, to be part of the motivation for BH's theory). However, at best it poses a challenge to explain why it might not be assertable, and since POP provides that explanation, the challenge is met.

8.7. Sorites

A good explanation of the Sorites should not just explain why Sorites arguments are unsound (as they surely are), they should explain why they seem sound. POP splits this task into three parts, and quickly disposes of two of them. (We will have to leave the fourth until the next section.) For simplicity, we will again use Tappenden's notation of writing $P(n)$ for 'A person with exactly n cents is poor'. (And we will assume, rashly, that poverty supervenes on net wealth.) Now consider the following three Sorites arguments.

- | | | |
|------|-----------------|--|
| (32) | Prem1. | $P(1)$ |
| | Prem2. | It is not the case that $P(1)$ and not $P(2)$ |
| | ... | |
| | Prem100,000,000 | It is not the case that $P(99,999,999)$ and not $P(100,000,000)$ |
| | C. | $P(100,000,000)$ |
| (33) | Prem1. | $P(1)$ |
| | Prem2. | If $P(1)$ then $P(2)$ |
| | ... | |
| | Prem100,000,000 | If $P(99,999,999)$ then $P(100,000,000)$ |
| | C. | $P(100,000,000)$ |
| (34) | Prem1. | $P(1)$ |
| | Prem2. | Either it is not the case that $P(1)$ or $P(2)$ |
| | ... | |
| | Prem100,000,000 | Either it is not the case that $P(99,999,999)$ or $P(100,000,000)$ |
| | C. | $P(100,000,000)$ |

The arguments are, I claim, arranged in decreasing order of intuitive plausibility of the premises. It feels absurd to deny, or even to decline to assent to, any of the premises in (32). One might resist some of the premises in (33), and in borderline cases the premises in (34) have very little persuasive force.¹⁴ Two conclusions immediately follow. We do not need to explain the intuitive force of (34), since it has no intuitive force. And our explanation of the intuitive force of (32) and (33) should not work equally well as an 'explanation' of the intuitive force of (34). This last condition is something of a worry for a few theories of vagueness, since without POP it is hard to see how, say, an epistemicist or a many-valued theorist could account for the difference between (32) and (34), their premises being trivially equivalent within their respective theories. POP is rather helpful here. According to POP, the objective semantic content of every premise in (32) is satisfied, but the objective pragmatic content of the premises in (34) concerning borderline cases are not satisfied. The problematic case is (33). Since, as I've said a few times, we don't really understand conditionals all that well, it is not easy to say much in detail about the

¹⁴ I am grateful here to conversations with Ted Sider.

objective pragmatic content of the premises in (33). The best we can say, and it isn't totally compelling, is that speakers accept the premises in (33) because they accept the premises in (32) and think (perhaps mistakenly) that they entail the premises in (32). For a quick introspective test of this, think of how you would react to someone who said they didn't see any reason to accept the premises in (33). Would you respond by pointing out to them that there can hardly be an n such that $P(n)$ but not $P(n+1)$, so if $P(n)$ then $P(n+1)$? If so, then you too think the main evidence for the premises in (33) is the premises in (32). This might be the correct explanation, but perhaps there is something else to say. To say it we need, however, to take one final detour through conditionals.

8.8. Modus Ponens, Epistemicism and Some Conclusions

Consider again this example, one of Vann McGee's counterexamples to modus ponens. (McGee 1985: 463; the numbering is not in the original)

I see what looks like a large fish writhing in a fisherman's net a ways off. I believe

(35) If that creature is a fish, then if it has lungs, it's a lungfish.

That, after all, is what one means by "lungfish". Yet, even though I believe the antecedent of this conditional, I do not conclude

(36) If that creature has lungs, it's a lungfish.

Lungfishes are rare, oddly shaped, and, to my knowledge, appear only in fresh water. It is more likely that, even though it does not look like one, the animal in the net is a porpoise.

What exactly should one make of this? We should all agree that (35) does in some sense follow from the meanings of the words it contains. I think that the sense is that its objective pragmatic content is satisfied, not that it is, say, true. After all, its antecedent is, we may suppose, true and its consequent apparently false, and I am somewhat more certain than conditionals with true antecedents and false consequents are false than that any semantic judgement I (or anyone else) would make about such a case is correct. But this flat-footed defence of the falsity of (35) does not explain its pragmatic acceptability. Presumably it is acceptable because the consequent follows, in some pragmatically salient sense, from the antecedent. But what is that sense?

One might be tempted to appeal to exportation here, but that is just a red herring. It is one thing to note that speakers typically assent to *If A and B, then C* just when they assent to *If A, then if B then C*. It is another thing altogether to explain why these patterns of assent sway together. If one knew that, one would know why (35) seems acceptable. If one does not, then one does not know why (35) is acceptable. So we must look deeper.

If POP is correct, then the explanation for the acceptability of (35), like that of (9), is that the objective pragmatic content of the consequent follows from the objective pragmatic content of the antecedent. So if we get clear on what the objective pragmatic content of *That is a fish*, we will might solve the puzzle, and learned something rather interesting about the pragmatic content of simple sentences.

One possibility can be quickly disposed of, that the objective pragmatic content of *That is a fish* is just that that is, indeed, a fish. If that were true we would expect the objective pragmatic content of (36) to follow from the fishiness of the object. As we have seen, it seems as if it does not. Even if (36) is true, and that might be the right thing to say about the example, we cannot assert it, and this does not seem to be merely because of ignorance.

Another possibility might be that the objective pragmatic content of *That is a fish* is that it is a metaphysically rigid fact that it is a fish. By ‘metaphysically rigid’ here I do not mean that it is true in all metaphysically possible worlds, just that it is true in all metaphysically nearby worlds. This *might* solve the problem. If there is some metaphysical stability to the fishiness of the object, then we might think that if *it* has lungs then *it* is indeed a lungfish, because *it* could not fail to be a fish, in a suitably strong sense of could. Whatever the merits of that suggestion, it seems implausible when applied to sentences outside conditionals. I can properly say *The Cubs were unlucky last century* without suggesting that their luck is anything other than a coincidence, so the pragmatic content of that sentence cannot be that there is anything metaphysically stable about its truth.

If metaphysical stability does not do the trick here, does epistemic stability do any better? Not if we mean by ‘epistemic stability’ that the speaker knows the sentence to be true. We noted above that assuming the objective pragmatic content of *S* is *The speaker knows that S* causes some difficulties for handling disjunctions. Perhaps, though, it might be suggested that the type of pragmatic content that gets embedded in conditionals is different to the type of content that gets embedded in disjunctions. This is ad hoc, but not altogether implausible. Still, it does not handle all the data about conditionals. As Richmond Thomason noted¹⁵, we can properly say sentences like *If the President is a spy, then we are all being successfully deceived*, where the consequent most assuredly does not follow from our *knowing* the antecedent.

Maybe a more subtle form of epistemic stability will work. Say that the objective pragmatic content of (a simple sentence) *S* is *It is humanly possible to know that S*. And assume that the right analysis of indicative conditionals is, broadly speaking, epistemic.¹⁶ So, roughly, *If A then B* is true if the epistemically nearest worlds where *A* is true are worlds where *B* is true. Now we have an explanation for why (35) seems acceptable. If it is knowable that the thing is a fish, and this has just been made salient, then we might suppose that the epistemically nearest worlds are all ones where it is a fish.¹⁷ Whether this is a plausible story at the end of the day will depend on how it links up with other facts about the behaviour of indicative conditionals. But for now we should just note that there is *some* chance that this explanation will go through. And, given the failures of other explanations to account for the plausibility of (35), this provides at least *some* evidence that part of the objective pragmatic content of *S* is that it is knowable that *S*. And that, in turn, provides some evidence that epistemicists about vagueness might be able to appeal to POP to explain why (4) is acceptable but (3) is not.

Given this, we *may* have one other explanation of the reason we are inclined to accept the premises in a Sorites argument where the main premises are conditionals, as in (33). If the objective pragmatic content of the antecedent is that the antecedent is knowable, then that certainly entails the

¹⁵ Cited in van Fraassen (1980)

¹⁶ As, for instance, the analyses of indicative conditionals in Stalnaker (1975), Davis (1979) and Weatherson (2001) are.

¹⁷ If one is wedded to a particular analyses of the indicative conditional, this move might seem a little quick. But if you think, as I do, that we do not know very much at all about the similarity metric relevant for assessing indicative conditionals, other than that it is sensitive to epistemic considerations and to salience, then we can take the little pattern of reasoning in the text to be a discovery about the nature of that metric.

truth of the consequent. So might that explain the acceptability of the conditional? This is possible, but it seems like an inferior explanation to the one provided at the end of the last section. The problem is that one would expect that the sentence would be acceptable if the objective pragmatic content of the antecedent implied the objective pragmatic content of the consequent, which it does not. And, of course, which it could not, unless somehow that objective pragmatic content was somehow immune to Sorites arguments. Which, obviously, it is not. So while getting clear on how indicative conditionals work might tell us something important about which aspects of content are compositional, it does not seem likely to explain why the main premises in arguments like (33) seem acceptable.

Let us take stock. We started with some puzzling data about reactions to sentences that look like logical truths, namely, that vagueness seems to threaten the law of excluded middle but not, in the first instance, the law of non-contradiction. And we proposed a theory, COP, that could explain this data, whether or not one accepted that all theorems of classical logic really are logical truths. Already, though, there was some difficulty about whether this explanation was consistent with an epistemic account of vagueness. We then noted some further data, concerning speakers' willingness to accept the law of excluded middle in some circumstances, but not accept some complex instances of the law of non-contradiction in others, that could not be explained by COP. So we posited a modification of COP, POP, that could explain even this new data. Unlike COP, POP did require us to assume that all truths of classical logic really are logical truths, but like COP, it was far from clear that POP was compatible with an epistemic account of vagueness. At this stage it looked as if we had something like an argument for a broadly supervaluational account of vagueness, or at least for some account of vagueness that preserves much of classical logic while accepting that vagueness is semantic indeterminacy. The considerations of this last section suggest we should temper our enthusiasm for that argument somewhat, because they show that epistemicism *might* be compatible with POP. Still, we already had several other arguments against epistemicism, so the fact that we don't get one more here shouldn't cause us to lose too much sleep. Of course, we also had several arguments against the only known *semantic* theory of vagueness, supervaluationism. Unless we can find a semantic theory of vagueness that avoids the problems that befell supervaluationism, then the fact that semantic theories integrate more smoothly with POP than epistemic theories will not count for much. The next chapter sets out just such a theory.