Indeterminacy as Indecision

The Woodbridge Lectures 2020

John MacFarlane

forthcoming in *The Journal of Philosophy*

Contents

Lecture I: Vagueness and Communication

I want to begin by reading a passage from a letter Gottlob Frege wrote to Giuseppe Peano in September 1896:

> A conditional definition of the sign for a concept decides only for some cases, not for all, whether an object falls under the concept or not; it does not therefore delimit the concept completely and sharply. But logic can only recognize sharply delimited concepts. Only under this presupposition can it set up precise laws. The fallacy known by the name of 'Acervus' rests on this, that words like 'heap' are treated as if they designated a sharply delimited concept whereas this is not the case. Just as it would be impossible for geometry to set up precise laws if it tried to recognize threads as lines and knots in threads as points, so logic must demand sharp limits of what it will recognize as a concept unless it wants to renounce all precision and certainty. Thus a sign for a concept whose content does not satisfy this requirement is to be regarded as meaningless from the logical point of view. It can be objected that such words are used thousands of times in the language of life. Yes; but our vernacular languages are also not made for conducting proofs. And it is precisely the defects that spring from this that have been my main reason for setting up a conceptual notation. The task of our vernacular languages is essentially fulfilled if people engaged in communication with one another connect the same thought, or approximately the same thought, with the same sentence. For this it is not at all necessary that the individual words should have a sense and meaning of their own, provided only that the whole sentence has a sense. Where inferences are to be drawn the case is different: for this it is essential that the same expression should occur in two sentences and should have exactly the same meaning in both cases. It must therefore have a meaning of its own, independent of the other parts of the sentence.[1]

The passage defends the uncompromising view of vagueness for which Frege is well known. If we seek rigor and want to apply logic,

[1] Gottlob Frege, *Philosophical and Mathematical Correspondence*, ed. Brian McGuinness, trans. Hans Kaal (Chicago: University of Chicago Press, 1980), pp. 114–15. I depart from Brian McGuinness's translation in translating 'Satz' as 'sentence' rather than 'proposition.'

Frege thinks, we must construct an artificial language in which every term has an independent meaning. In order to have an independent meaning (*Bedeutung*), a predicate must be completely defined: it must be settled whether it applies to any arbitrary object.[2] If one partially defines a mathematical function symbol—saying what values it takes on the natural numbers, for example, but saying nothing about the rationals—then the symbol lacks a Fregean *Bedeutung*. The vague terms of natural language, Frege thinks, suffer from the same problem. Although it may be settled that ripe red apples fall into the extension of 'red' while oranges do not, there are shades for which it is not settled whether they fall into the extension of 'red' or of 'orange.' If we treat these terms as if they are totally defined, then we face the sorites paradox: we are forced to classify each element in a series of slightly different shades as either falling into the extension of 'red' or not, and we find it impossible to justify giving two adjacent shades different classifications—yet we cannot classify them all the same way. We must, then, deny that vague terms have a *Bedeutung*, and this makes it inappropriate to apply logic to arguments in which they figure. It is not enough to avoid borderline cases: on Frege's view, logic cannot even be applied to an argument like 'Telly Savalas was a bald actor; therefore he was bald.'

This uncompromising attitude is a hard one to sustain. Frege did not succeed in his goal of creating a precise language adequate for pure mathematics, for his logicist theory harbored a latent contradiction (as Russell famously showed). We now make do with set theory, which arguably does not meet Frege's strictures, because its basic terms acquire meaning by partial definition, through their role in an axiomatic theory.[3] When we consider sciences like linguistics, economics, psychology, biology, and even arguably chemistry and physics, it becomes

[2]Gottlob Frege, *Posthumous Writings*, ed. Friedrich Kaulbach, Hans Hermes, and Friedrich Kambartel (Chicago: University of Chicago Press, 1979), p. 122.

[3]Solomon Feferman, "Why the Programs for New Axioms Need to Be Questioned," *Bulletin of Symbolic Logic*, VI, 4 (2000): 401–13 argues that the notion of set explicated in this theory is not determinate enough to settle the truth of important mathematical statements.

hopeless to maintain that our scientific language meets the standard of precision Frege set.[4] Since logic is evidently as useful in these sciences as it is in pure mathematics, Frege's dismissive attitude seems to go too far. We need a way of understanding vague discourse that allows it to be subject to logical criticism and control.

Over the last century, quite a bit of effort has been devoted to finding an alternative to Frege's dismissive attitude to vagueness. Two main approaches have emerged. The *semantic* approach seeks an alternative to classical semantics that is better suited for vague language. Instead of interpreting a language by assigning classical extensions to its predicates, we do something more complex—assigning fuzzy sets, for example, or sets of classical interpretations. In most cases the alternative semantics yields an alternative logic. The alternative logic and semantics is then brought to bear in explaining what goes wrong in sorites reasoning.

The *epistemic* approach, by contrast, argues that classical logic applies unproblematically to vague discourse, because vague words do have classical extensions, even if we cannot say what they are. On this view, there is a particular point at which adding a kernel of corn to some others yields a heap, a last second of one's youth, and so on. The reason we find it hard to believe these things, the epistemicist supposes, is that we are closet verificationists. We assume that if our words set up the sharp boundaries required by classical semantics, we would know where these boundaries lie. But why assume that? There are plenty of matters of fact of which we are ignorant. I have only a rough idea what the current air pressure is, for example, but there is definitely a fact of the matter about that. Why should we suppose that we are in a better position to know the semantic boundaries imposed by our use of 'red' and 'orange?'

The debates in this literature have focused largely on logical, semantic, and epistemological questions. What are so-called borderline

[4]Bertrand Russell, "Vagueness," *Australasian Journal of Philosophy and Psychology*, 1, 2 (1923): 84–92 embraced the conclusion that all scientific language is vague; for recent arguments that the fundamental notions of physics are vague, see Andrew Bacon, *Vagueness and Thought* (Oxford: Oxford University Press, 2018), pp. 233–5.

cases of vague predicates? Where does the fallacy lie in sorites reasoning? What sort of logical principles can we apply in the presence of vagueness? If our terms have classical extensions, why can't we know the locations of the boundaries they impose? Can it be vague whether something is a borderline case, and if so, what are the ramifications of this "higher-order vagueness?" How does vagueness relate to degree morphology, hedges, and other linguistic constructions?

These are interesting and important questions. But in these lectures, I want to focus on another question, to which much less attention has been given. Frege alludes to it briefly in the second part of the passage with which we began. After dismissing vague words as "meaningless from a logical point of view," he notes: "It can be objected that such words are used thousands of times in the language of life." If vague words lack an independent meaning, as Frege supposes, how do we manage to use vague language to communicate? To this question, Frege devotes just two sentences in response. And for the most part, the vast literature on vagueness has followed him in sidelining questions about ordinary communication in order to focus on questions of logic and formal semantics. I propose to do the opposite. We will dwell on the question of how we communicate with vague language, returning to the logical and semantic issues that have been central to the debate only at the end.

Here is the plan. In this lecture, I will try to get clearer about the problem vagueness poses for understanding communication, rejecting some attempts to dismiss it. I will argue that the problem is due not to the "fuzziness" of vague language, but to its contextual sensitivity. In fact, the problem is not specific to vagueness: it is a special case of a more general problem concerning felicitous contextual underspecification. I will argue that the traditional semantic and epistemic approaches to vagueness, which focus on the "fuzziness" of vague language, do not provide a good solution to this problem.

In the second lecture, we will look at some philosophers who have recognized the challenge posed by felicitous underspecification but gone in the wrong direction in meeting it. Their thought is that in

cases of felicitous underspecification, speakers express *clouds* of precise propositions. Instead of innovating in the theory of content, they propose that we innovate in the theory of speech acts. Although I do not think this approach is successful, the problems it faces point the way to a more adequate solution.

In the third lecture, I will present my own approach. I will argue that we need a new account of the contents of vague assertions, along the lines of the *plan expressivism* Allan Gibbard has advocated for normative language. On this view, the mental states we express in uttering vague sentences have a dual direction of fit: they jointly constrain the doxastic possibilities we recognize and our practical plans about how to draw boundaries. With this story in hand, we will reconsider some of the traditional topics connected with vagueness: bivalence, the sorites paradox, higher-order vagueness, and the nature of vague thought.

1.1 Frege's view

Let us have a look at those two intriguing sentences with which Frege dismisses the objection about vague communication:

> The task of our vernacular languages is essentially fulfilled if people engaged in communication with one another connect the same thought, or approximately the same thought, with the same sentence. For this it is not at all necessary that the individual words should have a sense and meaning of their own, provided only that the whole sentence has a sense.

Here, as elsewhere, Frege uses the word 'thought' to denote the content of a belief or assertion. He characterizes such a content as the mode of presentation of a truth value—and hence as a truth condition. So the model of communication Frege is assuming here is one we might call the

> Simple Model of Communication
> Communication succeeds just in case the hearer grasps the truth condition intended by the speaker.

In the simplest, non-figurative cases, the speaker who wants to communicate that *p* chooses a sentence whose truth condition is *p*; the hearer then recognizes that the sentence has this truth condition and infers that the speaker meant to communicate that *p*. Thus the hearer's epistemic achievement is mediated by a recognition of the sentence uttered by the speaker and speaker's and hearer's common knowledge of its sense.[5]

How does the hearer recognize the sense of the sentence? If it is a sentence in a language that meets Frege's strictures of rigor, each of its words (or other basic parts) can be associated with an independent meaning, and we can discover the sense of the whole sentence by composing the senses of the parts. In that case, there is an *algorithm* for determining the sense of the whole. But Frege seems to hold that the whole sentence can have a sense (in context) even if the parts do not all have independent meanings. In this case there is no mechanical procedure for determining the sense. Despite this, Frege seems to be saying, the hearer is often in a position to associate a definite truth condition with the sentence.

Frege does not give an example—or anything in the way of detail. But he may have had in mind this sort of case. A chemist says to her assistant,

(1) The tall test tube contains hydrofluoric acid.

The assistant is standing in front of a row of test tubes ranging in height from 25 cm (the tallest, on the left) to 10 cm (the shortest, on the right) (see Fig. 1.1).

[5]Compare David Lewis, "Index, Context, and Content," in *Philosophy and Grammar*, ed. Stig Kanger and Sven Öhman (Dordrecht: Reidel, 1980), 79–100, at pp. 80: "The foremost thing we do with words is to impart information, and this is how we do it. Suppose (1) that you do not know whether A or B or...; and (2) that I do know; and (3) that I want you to know; and (4) that no extraneous reasons much constrain my choice of words; and (5) that we both know that the conditions (1)–(5) obtain. Then I will be truthful and you will be trusting and thereby you will come to share my knowledge. I will find something to say that depends for its truth on whether A or B or ... and that I take to be true. I will say it and you will hear it. You, trusting

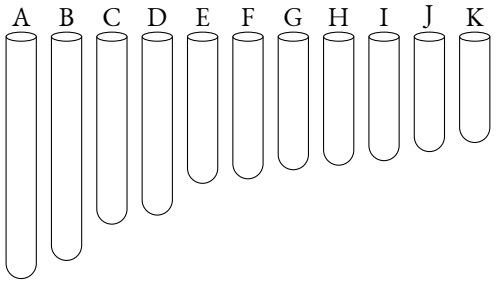

Figure 1.1: The test tubes.

It is easy, here, for the assistant to grasp the chemist's meaning: that test tube A contains hydrofluoric acid. The assistant can figure this out by drawing on two pieces of linguistic knowledge and one piece of nonlinguistic knowledge:

(a) 'The tall test tube' is only appropriate in a case where just one of the test tubes is presupposed to fall into the extension of 'tall.'
(b) If only one thing in a group falls into the extension of 'tall,' it is the tallest one.
(c) The tallest test tube is test tube A.

Thus, the sentence (1) can serve to get the chemist's message across even if 'tall' is not sharply defined. It is completely irrelevant to the inference here where exactly the line between the tall and the non-tall test tubes falls, or whether there is a line at all. Moreover, the inference here is a reliable one, capable of producing *knowledge* of what the speaker meant. It does not require past experience of a special use (like an idiom) or a guess about what is meant in context.[6]

I think that there are many cases like this, where a vague sentence can be used to communicate something precise. But can all cases of suc-

me to be willing and able to tell the truth, will then be in a position to infer whether A or B or ...".

[6] *Pace* Timothy Williamson, *Vagueness* (London: Routledge, 1994), p. 44.

cessful communication using vague language be handled in this way? That seems dubious. Suppose the chemist says,

(2) Some of the tall test tubes contain hydrofluoric acid.

It seems that the assistant can understand this utterance, too. Communication does not fail, as it would if the chemist said

(3) Some of the snurg test tubes contain hydrofluoric acid,

where 'snurg' is an unknown word, or if she said

(4) Those test tubes contain hydrofluoric acid,

without making any attempt to clarify which test tubes she is demonstrating.

But if successful communication requires the hearer to recognize the truth condition the speaker intends to communicate, it is puzzling how the utterance of (2) *could* be a case of successful communication. For recognizing the truth condition the chemist means to get across in uttering (2) would seem to require knowing where she puts the cutoff point for 'tall.' If she only regards test tubes 24 cm and taller as 'tall,' then what she is saying is true just in case some of the 24–25 cm test tubes contain hydrofluoric acid. She would be ruling out the possibility that only the 22 cm test tubes contain hydrofluoric acid. On the other hand, if she regards 21 cm test tubes as tall, she would not be ruling out this possibility. In the case of (1), we did not need to know where the cutoff point for 'tall' lies in order to recognize the truth condition the speaker meant to communicate. But in the case of (2), it seems, we do. It is difficult, then, for Frege's approach to make sense of successful communication with sentences like (2).[7]

[7] Williamson, *ibid.*, pp. 44–45 makes a similar complaint: "If I say to you 'A heap of beans will be placed on your doorstep at noon tomorrow,' you understand what I have said even though you have never encountered the sentence before and the context gives you no special clue as to its meaning."

1.2 Epistemicism

Frege's observation was that a vague sentence can be used to communicate something precise, when the context cooperates. We have seen that this approach does not go far enough: when the assistant utters (2), for example, there is no precise proposition that the chemist can recognize him as intending to get across. What shall we say about cases like this?

One possibility is that communication in such cases cannot be modeled as the hearer's recognition of a truth condition intended by the speaker. Perhaps the assistant's utterance of (2) is associated with a *cloud* of truth conditions rather than a single one. We will look at this option in Lecture II, but let us not give up so easily on the Simple Model of Communication.

Recall that for Frege, the sense of a word like 'tall' is a mode of presentation of a concept—that is, of a function from objects to truth values. Clearly, the chemist and the assistant do not share a mode of presentation for 'tall' of the form

(5) *x* is tall iff the height of *x* is greater than *n* mm.

But perhaps there is another mode of presentation that they *do* share: one that determines a cutoff point, but not in terms of millimeters. Perhaps they share a *vague* mode of presentation of the very same height concept that could also be picked out by a precise mode of presentation like (5). If so, then there *is* a truth condition the chemist intends to communicate with (2), and the assistant is in a position to recognize that truth condition, so the Simple Model of Communication works for this case.

Indeed, if the two parties *do* associate an independent *Bedeutung* with 'tall' via a vague mode of presentation, then (*pace* Frege) the assistant is in a position to work out compositionally what truth condition the chemist intends to communicate: namely, that some of the test tubes that have the property of being *tall* (or maybe *tall for a test tube in this lab*) contain hydrofluoric acid. Communication here is per-

fectly straightforward and does not require any fancy inferences about the context.

In rejecting the idea that vague words have a *Bedeutung*, then, Frege seems to be excluding the possibility that there could be a *vague* mode of presentation of a fully defined concept. It is not clear why. Two modes of presentation of the same thing can differ in many ways: why not in respect to vagueness? For example, consider three different modes of presentation of a certain Manhattan street corner:

(a) the NE corner of Amsterdam Avenue and W 106th Street.
(b) the corner with Mama's Pizzeria.
(c) the corner depicted in Fig. 1.2.

Figure 1.2: A Manhattan street corner. Photograph by Magnus Bäck, licensed under CC by 3.0, downloaded from Wikimedia Commons.

All three pick out the same corner, but they put their possessors in a different epistemic position with respect to that corner. (a) will be helpful in guiding you to the corner if you have an elementary grasp on Manhattan's street layout. (b) will help you if you have often walked past, even if you have not been keeping track of street numbers. (c) may not help guide you to the corner at all, but it will help you recognize it if you happen to pass by. In a similar way, there might be different modes of presentation of a height concept. A mode of presentation like (5) will

enable you to use a ruler to tell which objects are in the extension of *tall*. A vague mode of presentation will not do this, but that is no reason for thinking it is not a mode of presentation of the same concept. To be sure, the concept must somehow determine a cutoff point, if Frege's requirement that concepts be fully defined functions from objects to truth values is to be met. But nothing requires that this cutoff point be grasped in millimeters.

The basic point has nothing to do with vagueness. For example, in grasping the sense of 'at least as tall as the Empire State Building,' I grasp a fairly precise height concept. (For now, set aside issues about whether to include the antenna, what counts as the bottom of the building, and so on.) Yet I may have no idea how many *meters* tall something must be to fall under this concept. I do determine a threshold, but not under a mode of presentation of the form '*n* meters.'

You may feel that this is not enough to diffuse the objection. In the case of 'at least as tall as the Empire State Building,' there is something in my cognition that determines where the line is drawn. It does so indirectly, via comparison with an external object of whose height in feet I am ignorant, but we have something to point to in explaining why the line lies in one place rather than another. In the case of 'tall,' what can we point to?

Here epistemicists tend to just shrug. Somehow, they will say, the facts about our usage of 'tall' do determine a line, and we should not expect the relation between use and meaning to be transparent.[8] Critics of epistemicism have found this hard to accept,[9] but I am not going to contest the claim here. Instead, I will simply grant the epistemicist that a vague sense can determine a classical extension, in a way that is opaque to us.

This concession, by itself, is not enough to solve our puzzle about vagueness and communication. If communication is to succeed, on the

[8] *Ibid.*, pp. 205–12.

[9] See, for example, Vann McGee and Brian McLaughlin, "Review of Timothy Williamson, *Vagueness*," *Linguistics and Philosophy*, xxi, 2 (1998): 221–35; Vann McGee and Brian P. McLaughlin, "Logical Commitment and Semantic Indeterminacy: A Reply to Williamson," *Linguistics and Philosophy*, xxvii, 1 (2004): 123–36.

Simple Model, it is not enough for the speaker to express a determinate sense. The hearer also needs to *recognize* the sense the speaker intends. If this sense is determined in a non-transparent way by the speaker's usage, then it is very hard to see how the hearer is going to do this. We may suppose that the speaker and the hearer have somewhat different dispositions to apply 'tall.' We might then think that the vague senses they apply to 'tall' will also be different, and pick out different height concepts (with different thresholds). In that case communication on the Simple Model will not be possible. The hearer may know what sense *she* associates with 'tall,' but she will not be in a position to know what sense *the speaker* associates with 'tall.'

The standard epistemicist response to this sort of worry is to invoke anti-individualism.[10] Just as two people with very different beliefs about what it takes for someone to have arthritis can mean the same thing by 'arthritis' and share thoughts about arthritis, so people with different patterns of usage of 'tall' can mean the same thing by 'tall' and share thoughts about what is tall. The contents of their beliefs and intentions depend constitutively on the public meaning of the words they would use to express them. So our skepticism that the hearer's and speaker's dispositions to apply 'tall' match need not translate into skepticism about whether they attach the same sense to 'tall.'

One might worry that the "linguistic division of labor" that is present in Burge's 'arthritis' case is missing here. There are no experts or dictionaries who will tell us exactly how tall a test tube must be to fall into the extension of 'tall.' Nor is there any single particularly natural concept, in the continuum of height concepts corresponding to different cutoff points, that stands out as the most eligible referent of 'tall.' But what is most essential to anti-individualism is the idea that, by virtue of regarding themselves as meaning the same things by their words, speakers make it the case that their words should be interpreted in the same way. And we do have that here: speaker and hearer regard themselves as talking about the same height property, and they are not

[10] In the sense of Tyler Burge, "Individualism and the Mental," *Midwest Studies in Philosophy*, IV, 1 (1979): 73–122.

disposed to react to small discrepancies in classification by supposing that they simply attach different senses to the *word* 'tall.'[11] Because of this, it is the aggregate usage of the community, not the individual, that fixes the sense.

Let us grant *all* of this to the epistemicist: that vague modes of presentation can determine a Fregean *Bedeutung*, and that there can be shared contents even in cases where the manner in which meaning is determined by use is opaque to us. I think that the resulting position still faces a serious problem accounting for communication with vague language, which I will now describe.

1.3 Excursus: this and that

To see the problem, it will be helpful to start with a simpler case. Suppose the chemist says to her assistant,

(6) I want you to clean this test tube,

but does nothing to indicate which test tube she has in mind. I hope we can agree that communication has failed. The chemist may have *intended* to communicate something about a particular test tube, but because the assistant has no way of knowing which test tube, he is not in a position to recognize what she meant.

This is not to deny that the assistant has obtained some information about the chemist's intentions: that she wants him to clean a specific test tube, that it is proximal rather than distal (since 'this' was used rather than 'that'), and that she intends him to recognize which one she means. It is also not to deny that some of the speaker's aims in uttering (6) may be achieved: the assistant, afraid to ask for clarification, may simply wash all of the test tubes. Or the assistant may *guess* which test tube was intended, and the guess may be correct. In this case, the objectives the speaker wants to achieve by communicating may be met, *despite* the communicative failure. The notion of successful communi-

[11] *Ibid.*, sec. IIIc.

cation at issue here is that of success in an illocutionary act, not success in a perlocutionary act.[12]

I have been assuming so far that the chemist has a particular test tube in mind. If not, then there is a deeper failure: not only has communication failed, but the chemist has failed to mean anything at all.

For successful communication with a singular demonstrative like 'this,' then, it seems that (a) the speaker must intend to communicate something about a particular object, and (b) the hearer must recognize which object the speaker intends. As shorthand I will say that speaker and hearer must *coordinate on* an object.

Despite the strong ecological pressure to fill every unoccupied philosophical niche, I do not know of anyone who has been tempted to take an epistemicist position here, using moves parallel to the ones rehearsed above for 'tall.' Imagine the following dialogue:

> *Assistant:* Which test tube do you mean?
>
> *Chemist:* As I said, *this* one.
>
> *Assistant:* But I don't know which one you mean by 'this.' Do you mean test tube *A*, or test tube *B*, or...?
>
> *Chemist:* Well, *this test tube* is certainly one of the ones close to me, but I couldn't tell you which.
>
> *Assistant:* You said 'this test tube,' but you can't tell me which one you meant by 'this?'
>
> *Chemist:* Are you some kind of verificationist? It's a substantive question whether this test tube is test tube A or test tube B, and I don't need to be able to settle that question in order to refer to it using the term 'this test tube.'
>
> *Assistant:* But how can I understand you if I can't tell which test tube you meant?
>
> *Chemist:* It is enough for communication that we share a partial grasp of a socially individuated meaning. By 'this test tube' I mean this test tube; you understand me because you know what 'this test tube' means.

[12] J. L. Austin, *How to Do Things with Words* (Cambridge, MA: Harvard University Press, 1962), pp. 115–16.

The reason the appeal to anti-individualism is absurd here is that 'this' is highly context-sensitive. All uses of 'this' are subject to the constraint that the referent of 'this' should be proximal relative to the referent of 'that,' and one might invoke anti-individualism to argue that even users of 'this' who are not aware of this constraint are subject to it. But when it comes to the question which of several proximal objects the speaker is referring to with a particular use of 'this,' the relevant community is a community of one: the speaker has complete authority over this. If the hearer does not recognize which object the speaker intends to refer to, communication has failed. Appeals to the linguistic dispositions of other speakers, or to other uses of 'this' (for example, at the breakfast table earlier in the morning), are irrelevant.[13]

1.4 The contextual sensitivity of vague language

I want to suggest that the epistemicist's appeal to anti-individualism fails for vague words like 'tall' for a similar reason. To be clear: I do not think that we should think of 'tall' on the model of 'this,' as an ordinary context-sensitive expression. (I will be developing an alternative view in the next two lectures.) However, *if* the epistemicist is right that uses of 'tall' express vague modes of presentation that determine classical extensions, then 'tall' is like 'this' in the following respect. Its shared social meaning *constrains*, but does not determine, the mode of

[13] This is clear from thinking about how the speaker would react to being corrected. As Burge, "Individualism and the Mental," *op. cit.* notes in discussing his famous example, someone who claims to have arthritis in his thigh will react to being told that arthritis is a disease of the joints by saying, "I guess I was wrong—what I have in my thigh must not be arthritis." This shows that she intends to mean what others mean by 'arthritis.' She does *not* say, "You misunderstand me: what *I* mean by 'arthritis' is an inflammation of the joints *or* the bone." But when a hearer interprets a speaker as using 'this' to refer to a different object than the speaker took herself to be referring to, an accusation of misunderstanding is the only option. Kaplan's much-discussed Carnap/Agnew example ("Dthat," in *Syntax and Semantics, Vol. 9: Pragmatics*, ed. Peter Cole (Academic Press, 1978), 221–43, at p. 239) does not show otherwise. In the scenario Kaplan describes, he intends to use 'that' to refer *both* to the picture he's pointing to and to a picture of Carnap. His false belief that these pictures are one and the same makes it difficult to say what he intends to refer to, but we do not get a clear judgment that he is referring to a different object than he intends.

presentation it can be used to express in a context. The speaker can use it in one way or in another, and for communication to be successful (at least on the Simple Model) the hearer must recognize which way that is. Anti-individualism can help secure a common standing meaning for 'tall,' but it cannot help with this second task: coordinating on one of the multiple possible uses permitted in the context by the shared standing meaning.

The problem, in short, is this. Although we should concede that a vague sentence like (2) can be used to communicate a vague thought, there is not just one vague thought one might seek to communicate with that sentence: there are many. If this is right, then for communication to succeed, on the Simple Model, the hearer needs to be able to recognize which of these is the speaker's intended meaning. And anti-individualism cannot help with this, because *all* of the candidate speaker meanings are permitted by the shared standing meaning of the sentence.

One way to resist this argument would be to claim that there is just one sense that a speaker can express using 'tall.' I do not know anyone who would go this far. "Semantic minimalists" who deny that 'tall' is a context-sensitive word, as a matter of its semantics, will happily concede that speakers can use sentences containing 'tall' to communicate thoughts involving a great variety of height properties.[14] I take it as uncontroversial that, in a particular context, a speaker might say either

(7) Those children are not tall,

meaning that they are not tall for a human, or

(8) Those children are tall,

meaning that they are tall for American third-graders. Whether the multiplicity of possible speaker meanings is due to the semantic context

[14] Herman Cappelen and Ernie Lepore, *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism* (Oxford: Blackwell, 2005) recognize that their semantic minimalism would not be plausible without "speech act pluralism."

sensitivity of 'tall' or is entirely pragmatic is irrelevant to the question at issue, which is how the hearer can recognize what the speaker means.

A more plausible line of resistance would be to claim that in any given context, there is a fairly restricted number of things a speaker might mean by a sentence like (2) or (8), and that in normal cases it is possible for the hearer to work out which of these is the intended message. For example, if the options in the case of (8) are

(a) Those children are tall for American third-graders.
(b) Those children are tall for American students.
(c) Those children are tall for an American human.
(d) Those children are tall for a human.

then it may be pretty obvious that the speaker intends (a). Of course, speaker and hearer may not agree on exactly how tall someone has to be to be tall for an American third-grader. But if we can be anti-individualists about 'tall for,' then the only difficulty lies in selecting the reference class. And perhaps that is no harder a task than coordinating on the referent of 'this.'

But this response oversimplifies the problem, for three basic reasons. I will summarize them here, then discuss them in more detail:

1. The possibilities for the implicit reference class are not so neatly bounded. The problem of coordinating on a reference class is thus significantly harder than the problem of coordinating on a referent for 'this.'

2. Indeed, it is not clear that the speaker normally has a particular reference class in mind.

3. Even when the speaker does have a particular reference class in mind, this is not enough to determine the speaker's meaning. Even relative to the same reference class, the speaker often has a choice between incompatible senses of 'tall' that draw height boundaries (or fuzzy height boundaries, if you wish) in different places.

To appreciate the first problem, one has only to see that in addition to option (a) above, the following would all be reasonable implicit reference classes to use with 'tall' in (8):

(e) Those children are tall for third-graders in this country.
(f) Those children are tall for third-graders in this city.
(g) Those children are tall for third-graders in this school.
(h) Those children are tall for third-graders in this class.
(i) Those children are tall for third-graders in their socioeconomic class.
(j) Those children are tall for the third-graders we know.

and so on. It is not going to be obvious to most hearers which of these the speaker intends.

One might suppose that we face a very similar problem with 'this.' Even if a speaker is pointing, there are indefinitely many objects through which the ray extending from their fingertip passes. That couch, yes—but also that cushion, and that ten-inch square of fabric on the cushion, and the stuffing just beyond the fabric, and the wall socket behind the couch, and so on. However, it is much easier to recognize which object speakers intend to demonstrate in using 'this' than it is to recognize which reference class they have in mind. If one points at a couch and says

(9) This looks expensive

one will normally be assumed to be referring to the couch. If one wants to refer to the fabric, or a square of fabric, or the wall socket behind the couch, one will have to do something extra to indicate this. None of the candidates (e)–(j) for reference classes enjoys this kind of "default" status.

Indeed, it seems doubtful that in most cases speakers have a determinate enough intention to pin down one of these candidates and exclude all the others. If you ask someone who utters (8)

(10) Do you mean tall for a third-grader in this school, or in this class, or in this city?

you are likely to get the answer, "I don't know, I didn't really have any of those specifically in mind." Even if you do get a definite answer, how confident will you be that it reports a pre-existing intention, and not a new decision made on the spot in response to your question? Here the analogy with 'this' breaks down completely. It seems definitely defective to use 'this' without having a determinate object in mind, but using 'tall' without having chosen between, say, (f) and (g) does not seem infelicitous in the same way.

What is more, even if speakers *did* have determinate reference classes in mind, and even if hearers *could* recognize these, this would not exhaust the contextual sensitivity of gradable adjectives.[15] As Delia Graff Fara observes, 'old for a dog' can mean different things:

> Suppose that Fido is fourteen years old and Rover is twenty years old. Someone who says that Rover is old for a dog may be making a remark about his extreme longevity, while someone who says that Fido is old for a dog may be merely remarking that he is in his old age.[16]

She proposes that, in addition to a reference class, 'old' is also sensitive to a contextually varying *norm* that determines a "typical" age within the selected comparison class. To be *old for an F* is to be significantly older than a typical *F*:

> Rover, the twenty-year-old, has significantly more age than is the norm for a dog to attain; while Fido, the fourteen-year-old, has significantly more age than some different kind of norm for a dog, one that's much harder to articulate, but which perhaps concerns the peak age of good health.[17]

[15] In addition to the works cited below, the point is made in Christopher Kennedy, "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives," *Linguistics and Philosophy*, xxx, 1 (2007): 1–45 [sec. 2]; Mark Richard, *When Truth Gives Out* (Oxford: Oxford University Press, 2008) [chap. 4]; and John MacFarlane, "Vagueness as Indecision," *Aristotelian Society Supplementary Volume*, xc, 1 (2016): 255–83.

[16] Delia Graff Fara, "Shifting Sands: An Interest-Relative Theory of Vagueness," *Philosophical Topics*, xxviii, 1 (2000): 45–81, at p. 66.

[17] *Ibid.*

For a content to be communicated using a gradable adjective, then, speaker and hearer would have to coordinate not just on a reference class but on a norm of typicality. Consider again our example of the chemist and her assistant, with the test tubes arrayed as depicted in Fig. 1.1. Which test tubes are tall for a test tube in this group? A, B, C, and D? Or just A and B? It seems to me that 'tall for a test tube in this group' can be used in either way.

The idea that the contextual sensitivity of gradable adjectives is limited to the provision of a reference class is even less plausible when we note, with Keith DeRose, that there are uses of gradable adjectives that do not seem to have any connection to a reference class at all:

> Consider a movie director setting up the background for a key scene. 'I need something tall over there on the left, to balance the shot—maybe a tree, or a streetlight or something. Get me something tall!' When asked what she means by 'tall,' the natural answer for her to give might well make no reference at all to any comparison class, but might rather simply cite what we may call a 'tape-measure' standard: 'What I mean is something about 14 to 16 feet tall.' It might not only be natural not to cite any reference class, but it might be hard to say what comparison class could possibly be instead cited. The director has mentioned trees and streetlights, but she's not looking for things that are tall for a tree or tall for a streetlight; in fact, what she seeks might well be quite short for a tree or for a streetlight. Is there some implicit reference class that the director hasn't mentioned?[18]

In addition to the threshold-related context sensitivity we have been discussing so far, the underlying *ranking* of objects as more or less tall would seem to vary with context. For example, in one context we may decide to count architectural spires in determining a building's height, in another not. (In 1929, the Chrysler Building's claim to be the world's tallest building turned on the answer to this question.[19]) And of course contextual variation in the ranking is much greater for

[18] Keith DeRose, "Gradable Adjectives: A Defence of Pluralism," *Australasian Journal of Philosophy*, LXXXVI, 1 (2008): 141–60, at p. 149.

[19] In what was then called the "Race into the Sky," the Chrysler Building and the Bank of Manhattan Building (40 Wall Street) competed to be the tallest. The Chrysler Building triumphed by including a spire which was constructed in secret and raised

"multidimensional" predicates like 'smart' or 'large.' In determining whether one thing is larger than another, one has to settle how to take into account differences in height, width, and depth, and there is not a single privileged way of doing this.

To sum up: *if* 'tall' has a classical extension, as the epistemicist supposes, and the Simple Model of Communication is correct, then we cannot explain how we manage to communicate using vague words. For on these assumptions, successful communication would require the speaker and hearer to coordinate on one of the many truth conditions that the sentence can be used to express in that context. Not only is this an insoluble problem on the hearer's side, but it is dubious that speakers even have determinate intentions for the hearers to discern.

1.5 Felicitous underspecification

Most investigations of vague predicates have focused on what Crispin Wright[20] called their "tolerance." Intuition rebels at the idea that tiny differences in height can affect whether something should be classed as tall, and that is what makes the sorites paradox plausible. The problem posed by vagueness is often conceptualized as the problem of rejecting the paradoxical tolerance intuition while explaining why it seems compelling.

The problem I have been trying to isolate, though, has no immediate connection to tolerance. I have said nothing about the sorites or borderline cases. And it should be clear that the usual semantic alternatives that have been proposed to make sense of these phenomena do not help with our problem. A move to a fuzzy semantics, for example—assigning degrees of truth to sentences, on a continuum from 0 to 1—

at the last minute. Architects for the Bank of Manhattan Building argued that their building should count as taller, on the grounds that it contained a higher usable floor (Wikipedia contributors, "40 Wall Street," *Wikipedia, The Free Encyclopedia*, 2020, https://en.wikipedia.org/w/index.php?title=40_Wall_Street&oldid=941342831, accessed February 28, 2020).

[20] "Language Mastery and the Sorites Paradox," in *Truth and Meaning*, ed. Gareth Evans and John McDowell (Oxford: Oxford University Press, 1976), 223–47, at p. 229.

can be seen as accommodating the intuition of tolerance, since now we can say that small changes in the object go with small changes in the degree to which a sentence is true. But fuzzy semantic values just make the problem of communication worse. If it is hard to coordinate on a classical extension (a function from objects to classical truth values), it is going to be even harder to coordinate on a fuzzy extension (a function from objects to real-numbered degrees of truth). In a domain of two objects, for example, there are only four possible classical extensions, but uncountably many possible fuzzy extensions. At any rate, moving to fuzzy semantic values does nothing to solve the problems we have been considering.

Similarly, keeping classical extensions but adopting an extremely fine-grained view of vague propositions, as Andrew Bacon has advocated, can make the coordination problem even more intractable. On Bacon's view, there is a vague proposition corresponding to every "evidential role"—every way of assigning probabilities to maximally consistent precise propositions.[21] Suppose that the particular proposition Alison means to express using 'Harry is bald' has an evidential role that assigns probability 0.733 conditional on Harry having 233 hairs (and other relevant precise facts). How on earth is a hearer supposed to recognize that the speaker intends this one, and not one that assigns the conditional probability 0.698? An appeal to anti-individualism seems the only hope, but as argued above, this is not enough, given the contextual fluidity of vague terms.

Our problem, then, is not due to the tolerance, "fuzziness," or sorites-susceptibility of vague language. It is a more general problem, which can arise even in cases where these features are absent: the problem of making sense of what Jeff King has called "felicitous underspecification."[22] Felicitous underspecification occurs when failure to coordinate on a determinate "supplement" for a context-sensitive term does not seem to be required for communicative success, in the way that it

[21] *Vagueness and Thought*, *op. cit.*, p. 118.

[22] Jeffrey C. King, "Strong Contextual Felicity and Felicitous Underspecification," *Philosophy and Phenomenological Research*, xcvii, 3 (2018): 631–57.

is for (6). King gives a number of examples.[23] For example, looking at a group of surfers down the beach, one might say

(11) Those guys are good.

Which plurality of surfers in that area is being referred to? The speaker may have no precise idea. Nor does the hearer have to answer this question in order to understand the speaker's assertion. Note that in this case there need not be a spectrum of possibilities: there might just be a few people around the edges whose inclusion in the group is unsettled. So this case need not involve "tolerance" or the other hallmarks of vagueness.

Planning a wedding in California, one might say

(12) Let's go with a local firm.

Local to the town, the county, the larger area? Nothing about the speaker's intention may decide that, though it might be clear that these are the relevant options. Similarly, 'Sophie's skis' can mean the skis belonging to Sophie, the skis Sophie made, or the skis Sophie is using. But in a case where Sophie is using her own skis, which she made, a speaker might use 'Sophie's skis' without distinguishing between these alternatives.

What these examples suggest is that it is not essential for successful communication that speakers and hearers coordinate on all the contextual supplementations that would be necessary to get to a proposition that is determinately true or false. If that is right, we must give up the Simple Model of Communication.[24]

[23] Examples of this sort are not new, of course. See, for example, François Recanati, *Literal Meaning* (Cambridge: Cambridge University Press, 2003) on 'here' and 'now' or Cappelen and Lepore, *Insensitive Semantics*, *op. cit.* on 'ready.'

[24] King suggests, plausibly, that what is common to these cases is that "singling out a unique semantic value for the expression is not crucial to the communicative aims of the speakers" ("Strong Contextual Felicity and Felicitous Underspecification," *op. cit.*, p. 651). If one took "communicative aims" to include perlocutionary aims (see p. 15, above), then one might try to explain these cases as illocutionary failures (and hence as compatible with the Simple Model). If all I care about is that you break some glass, and I say "break that glass!" without indicating which one I mean, then

1.6 Why not diagonalize?

Described abstractly, these cases of felicitous underspecification are cases in which it is not common ground what the semantic value of a context-sensitive expression is in the context. Readers of Robert Stalnaker's paper "Assertion"[25] may therefore wonder whether the interpretive strategy outlined there, commonly called "diagonalization," can help us here too.

Stalnaker is assuming what we have called the Simple Model of Communication: in asserting a proposition, the speaker intends the hearer to recognize which proposition she is asserting. If the assertion is accepted, this proposition will be added to the conversational *common ground*—the stock of propositions that are taken for granted in the subsequent conversation. As Stalnaker observes, this picture requires that it be common ground which proposition the speaker has asserted. Otherwise, it would not be common ground how the common ground was supposed to be updated after the assertion was accepted. If we represent the common ground as a *context set*—the set of worlds not ruled out by propositions accepted in the common ground—then this constraint can be articulated as

> Uniformity
> The same proposition is expressed relative to each possible world in the context set.[26]

Stalnaker argues that our interpretation of what a speaker means to assert is often guided by the requirement of Uniformity. Suppose we hear a woman talking in the next room, and I say

---

my illocutionary act has failed, but my perlocutionary act may succeed, as long as you respond by breaking a glass. As I read King, though, his point is that coordination on a unique value is not required for the *illocutionary* aims of the speakers, and perfect uptake does not require such coordination. This position is a decisive rejection of the Simple Model.

[25] "Assertion," in *Syntax and Semantics, Vol. 9: Pragmatics*, ed. P. Cole (New York: Academic Press, 1978); reprinted in Robert Stalnaker, *Context and Content: Essays on Intentionality in Speech and Thought* (Oxford: Oxford University Press, 1999), pp. 78–95.

[26] Stalnaker, *Context and Content*, *op. cit.*, p. 88.

(13) That is either Zsa Zsa Gabor or Elizabeth Anscombe.

The common ground leaves it open that the woman is Zsa Zsa Gabor, and also that she is Elizabeth Anscombe. If she is Zsa Zsa Gabor, then the referent of 'that' in context is Gabor; if she is Elizabeth Anscombe, then the referent of 'that' in context is Anscombe. So, if (13) were being used to assert a singular proposition—if the demonstrative 'that' in (13) rigidly denoted a particular woman—then Uniformity would be violated. It would not be common ground which proposition was being asserted. In order to respect Uniformity, then, we interpret the speaker as asserting a different proposition: the proposition that the woman who is talking in the next room (whoever she might be) is either Zsa Zsa Gabor or Elizabeth Anscombe. This proposition, called the *diagonal*, is true at a world w just in case the proposition that (13) expresses if w is the actual world is true at w (see Fig. 1.3).

	g	*a*	*c*
g	T	*T*	*T*
a	*T*	T	*T*
c	*F*	*F*	F

$\Rightarrow$

g	*a*	*c*
T	T	F

Figure 1.3: Diagonalization. Here g is a world where Zsa Zsa Gabor is the woman speaking, a is world where Elizabeth Anscombe is the woman speaking, and c is a world where Tricia Cox Nixon is the woman speaking. In the "propositional concept" on the left, the rows represent different possibilities for the actual world, while the columns represent the world of evaluation. On the right is the diagonal proposition derived from this propositional concept.

It might seem that this strategy can be applied straightforwardly to all of our cases of contextual underspecification. For in all of these cases, we lack common knowledge of the extension of a contextually sensitive element, and hence of which proposition is literally expressed

by the sentence in question. In our test tube example, for instance, the diagonal of (2) would be

(14) Some of the test tubes that surpass the threshold for being tall (for a test tube) in this context contain hydrofluoric acid.

Since we do not know where the threshold lies, this proposition does not tell us anything definite about the height of the test tubes containing hydrofluoric acid. But it does tell us something nontrivial: the height must be above the threshold. If we later learned which test tubes contained hydrofluoric acid, we could use this information to infer the location of the threshold, and if we later learned something about the location of the threshold, we could infer which test tubes contained hydrofluoric acid.

Diagonalization, if it worked here, would give us a way to keep the Simple Model of Communication, since it offers a way for the speaker and hearer to coordinate on truth conditions when they cannot coordinate on a completion for a contextual parameter. But one might worry that this trick works *too* well. For, if we can diagonalize here, we should be able to diagonalize our way out of paradigm cases of communicative failure, like example (6), where the chemist uses 'this test tube' without a particular test tube in mind, and without doing anything to indicate to the hearer which test tube she intends. Considered abstractly, this case looks a lot like Stalnaker's Gabor case. There is not common knowledge about the referent of 'this test tube' in the context. So, one might think, the worlds in the context set do not all agree about what the referent is, and we can diagonalize, coordinating on the proposition

(15) I want you to clean the test tube that is the referent of 'this' in this context.

But if we can diagonalize in this case, then it should be perfectly felicitous to assert (6) without doing anything to indicate what the referent is, and without settling oneself on a particular referent. This should be just as felicitous as asserting (13) when one does not know whether

'that woman' refers to Gabor or Anscombe or Cox. But it is not. Communication fails in the case of (6) in a way it does not in the case of (13). The cases must be different, then, in a way that blocks the diagonalization strategy for (6).

Here is where I think the difference lies. In the case of (13), speaker and hearer know how the referent of 'that woman' depends on worldly facts. They mutually know that it refers to the woman talking in the next room, whoever that may be. If Gabor is talking in the next room, then it is Gabor; if Anscombe, Anscombe; and so on. It is this mutual knowledge that allows them to fill in the squares of the matrix in Fig. 1.3 and to coordinate on the diagonal. But this is just what they lack in the case of (6). They know that if the speaker's intentions and public demonstrations do not fix a referent for 'this test tube,' nothing does. So if there are worlds in the context set in which the speaker has no determinate intention at all—as there will be in this case—there will be no way to fill in the squares of the matrix on the rows corresponding to those worlds, and the diagonal will not be well defined.

I want to suggest that (2) and the other examples of felicitous underspecification are resistant to diagonalization for the same basic reason as (6)—even though, unlike (6), they are felicitous. In all of these cases, we suppose that if the speaker's intentions and public demonstrations do not fix a certain contextual completion, nothing does. For example, if the speaker's intentions do not settle whether Zeke is a member of the group of surfers denoted by 'those guys' in (11), then nothing does. And if the speaker's intentions in (2) do not settle on one of the many contextually permissible ways of drawing a threshold for tallness for test tubes, then nothing does.

Talking of "the threshold of the context" is just as misguided as talking of "the referent of 'those guys' of the context" or "the referent of 'this test tube' of the context." Such phrases are not part of ordinary discourse; they belong to the language of formal semantics, and our familiarity with that language may make this talk seem less strange than it should. So it might be helpful here to consider the account of deictic pronouns in Irena Heim and Angelika Kratzer's widely used

semantics text.[27] Here deictic pronouns like 'she' are taken to be free variables. They are thus given their referents by the variable assignment of the context, g_c. Offloading the determination of reference for deictic pronouns into a contextually provided assignment function makes sense, since (aside from a gender presupposition) the meaning of 'she' determines nothing about the referent. But the variable assignment g_c is not a determinant of the referent of pronouns that functions independently of the speaker's intentions. It is, rather, just a way of representing the upshot of the facts about the speaker's intentions and any other relevant factors. Nor do Heim and Kratzer assume that the context always provides a value to a given variable. They say "let 'g_c' stand for the variable assignment determined by c (*if any*),"[28] and they add that a context is *appropriate for* a sentence only if it determines values for all the free variables in the sentence. It is only if we forget about the fact that, when a speaker lacks the requisite intentions, a context will fail to determine a variable assignment, that we will be tempted to think we can always diagonalize.

What I have been arguing is that talk of "the threshold determined by the context," and the use of the diagonalization strategy to interpret an assertion of (2), presupposes that it is taken for granted among speaker and hearer that the context does determine a threshold. Perhaps this is taken for granted in a small number of conversations taking place in Oxford. But in general, we do not think there is a fact of the matter about where the contextually determined threshold for 'tall test tube' lies. For this reason, we cannot diagonalize ourselves out of the difficulty we have posed for the Simple Model of Communication. In the cases that interest us, the presupposition that must be met in order for diagonalization to be possible is not met.

Let me emphasize that what matters here is what the conversational participants presuppose, not what is really the case. If the participants presuppose that facts about their context do determine a unique threshold for 'tall test tube,' then they can diagonalize and coordinate on an

[27] *Semantics in Generative Grammar* (Oxford: Blackwell, 1998), sec. 9.1.2.
[28] *Ibid.*, p. 243, emphasis added.

asserted content. If they do not presuppose this, then they cannot. So, even if the epistemicists were right about the metaphysics, and the context did somehow determine a unique threshold, this could help to explain vague communication in the (normal) cases where this is not presupposed. Cases like (2) would still present us with a thorny problem for the Simple Model of Communication, and the epistemicists' "actual threshold" would be an idle wheel in explaining the use of vague language.

1.7 Conclusion

I have argued that making sense of the use of vague language in communication is a genuine problem, distinct from the usual logical and semantic preoccupations of the literature on vagueness. The problem is that if communication requires recognition of the proposition a speaker intended to assert (the Standard Model of Communication), it is hard to see how communication with vague words can be successful when it is.

This is not, in the first instance, a problem in semantics, but a problem in pragmatics. It arises at the level of speaker meaning, not word meaning. It is a special case of the more general problem of explaining what King has called felicitous underdetermination. Once we appreciate the generality of the problem, we can see that it does not arise from the tolerance, "fuzziness," or sorites-susceptibility of vague language, and cannot be solved by remedies directed at these features (such as fuzzy truth values). Nor can it be dissolved by the usual maneuvers available to epistemicists.

In the next two lectures we will explore how the more general problem might be solved by rethinking the Standard Model of Communication.

Lecture II: Seeing Through the Clouds

In Lecture I, I argued that vagueness poses a problem for what I called the Simple Model of Communication. According to the Simple Model, communicative success consists in the hearer's grasp of the truth condition intended by the speaker. The problem is this: even if we allow that the truth condition intended by the speaker may be *vague*, there are so many candidate truth conditions that a speaker might intend in uttering a vague sentence that the hearer's epistemic task—the task of recognizing which truth condition is intended—is intractable. Indeed, in most cases speakers themselves do not seem to have a particular candidate in mind. The Simple Model predicts that these shortcomings should result in communicative failure, on par with what happens when one says

(1) That thing is interesting

without giving the hearer any clue as to which of several salient objects one intends to refer to, and without deciding oneself. But they do not. One can felicitously assert

(2) Some of the tall test tubes contain hydrofluoric acid

without deciding whether one is using 'tall' in a way that excludes test tube C (see Fig. 2.1, repeated from Lecture I). And others can count as understanding such an assertion. The Simple Model cannot make sense of this.

I argued that the problem vagueness poses for communication is just a special case of a more general phenomenon Jeff King has called "felicitous underspecification." The very same issue arises in cases that are not usually thought of as cases of vagueness. To use one of King's examples, the plural demonstrative in

(3) Those guys are good

refers to a group of people, but the speaker may not have a definite intention that settles whether

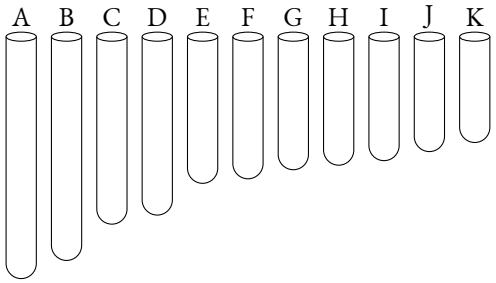

Figure 2.1: The test tubes.

(a) *those guys* = Abe, Bob, Cindy, Maria, Zeke, or
(b) *those guys* = Abe, Bob, Cindy, Maria, Sid.

Despite this, the assertion of (3) can be felicitous, and understanding it does not seem to require deciding between (a) and (b).

It is natural to think that the problem with the Simple Model is its assumption that the content asserted is a single proposition. If we gave up that assumption, we could say that in uttering (3) the speaker puts forth a *cloud* of propositions (perhaps, in this case, just two). Our task in this lecture is to explore this idea, which has been endorsed recently by a number of philosophers.[1] I will call the general idea, which different thinkers develop in different ways, "the cloudy picture."

In thinking about the cloudy picture and its alternatives, it is helpful to factor the Simple Model into two separable commitments. First,

[1]David Braun and Theodore Sider, "Vague, so Untrue," *Noûs*, xli, 2 (2007): 133–56; Ray Buchanan, "A Puzzle about Meaning and Communication," *Noûs*, xliv, 2 (2010): 340–71; Kai von Fintel and Anthony S. Gillies, "'Might' Made Right," in *Epistemic Modality*, ed. Andy Egan and Brian Weatherson (Oxford: Oxford University Press, 2011); Jeffrey C. King, "The Metasemantics of Contextual Sensitivity," in *Metasemantics*, ed. Alexis Burgess and Brett Sherman (Oxford: Oxford University Press, 2014), 97–118 [p. 106]; Justin Khoo and Joshua Knobe, "Moral Disagreement and Moral Semantics," *Noûs*, lii, 1 (2018): 109–43; Jussi Suikkanen, "Contextualism, Moral Disagreement, and Proposition Clouds," ed. Russ Shafer-Landau, *Oxford Studies in Metaethics*, xiv (2019): 47–69; Michael Caie, "Semantic Indecision," *Philosophical Perspectives*, xxxii, 1 (December 2018): 108–43.

there is a view about communication that abstracts from the nature of the contents communicated:

Classical Pragmatics

1. The content of an assertion is a (single) proposition.
2. Uptake consists in recognizing the proposition asserted.
3. If the assertion is accepted, its content is added to the conversational common ground.

Then, there is a view about these contents:

Classical Contents
Contents are ways the world could be (truth conditions).

The Simple Model results from the conjunction of Classical Pragmatics and Classical Contents.

The way the cloudy picture rejects the Simple Model is by rejecting Classical Pragmatics. Defending the cloudy picture, then, requires giving us a new story about assertion and uptake that can replace Classical Pragmatics. In what follows, I will argue that none of the proponents of the cloudy picture have succeeded in giving us a workable alternative to Classical Pragmatics, and that the difficulties they face can only be surmounted by rejecting Classical Contents. Thus, we can continue to think of communication as consisting of the hearer's grasp of the content intended by the speaker, but we must stop thinking of this content as a truth condition or "way the world might be." The question of how, exactly, we are to think of the content will be left for Lecture III.

2.1 Motivating clouds

We will look at three versions of the cloudy picture. Although they differ in significant ways and focus on different bits of language, they are motivated by similar considerations. The authors present cases in which the speaker's intentions do not single out one of a number of possible assertible contents. They then propose that, instead of taking the speaker to be expressing a single proposition, we take her to be expressing a whole cloud of them.

David Braun and Theodore Sider[2] focus on vague expressions. They assume that the world itself is not vague, and that propositional truth is not relative to anything other than the state of the world. It follows that every proposition is either determinately true or determinately false at a given world. That is, propositions are precise. Vague sentences, then, do not express propositions. Even in a specific context, a vague sentence expresses a cloud of propositions. For example, the sentence 'Bill is tall' might express a cloud containing the proposition that Bill exceeds 196 cm in height, the proposition that Bill exceeds 197 cm in height, and many others. As Braun and Sider put it:

> There is typically a cloud of propositions in the neighborhood of a sentence uttered by a vague speaker. Vagueness prevents the speaker from singling out one of these propositions uniquely, but does not banish the cloud.[3]

Ray Buchanan's[4] appeal to clouds of propositions is motivated not by lexical vagueness, but by cases of felicitous underspecification. For example, suppose Chet asserts

(4) Every beer is in the bucket.

Which proposition he has asserted depends on the quantifier domain, which is left implicit in (4). Does Chet mean

(5) Every beer we bought at the bodega is in the bucket,
(6) Every beer we will serve at the party is in the bucket,
(7) Every beer for our guests is in the bucket, or
(8) Every beer in the apartment is in the bucket?

Buchanan suggests, plausibly, that Chet's intentions may not distinguish between these options. If asked to clarify what he meant, he might respond with one of these propositions. But there is no reason to think that he had this one, rather than the others, in mind when he uttered (4).

[2] "Vague, so Untrue," *op. cit.*
[3] Braun and Sider, "Vague, so Untrue," *op. cit.*, p. 135.
[4] "A Puzzle about Meaning and Communication," *op. cit.*

> The fact that the speaker might, as it were, "fall back" on any one, or more, of the foregoing candidates, suggests that no single such candidate, or set of candidates, perfectly capture his communicative intentions in uttering (4). Chet's communicative intentions, such as they are, exhibit a certain kind of *generality* and *indifference* that precludes us from identifying any one of the candidate propositions as *the one he meant.*[5]

Buchanan concludes that

> The object of Chet's communicative intentions is not a proposition, but rather a property of propositions. ... Chet's utterance is, in some sense, "associated" with many non-truth conditionally equivalent propositions—namely, those propositions that are of the intended (restricted) type.[6]

Kai von Fintel and Anthony Gillies[7] also motivate the cloudy picture using a case of felicitous underspecification, which they attribute to Chris Potts:

> Billy meets Alex at a conference, and asks her:
>
> (9) Where are you from?
>
> That question is supposed, given a context, to partition answer-space according to how low-level in that context Billy wants his details about Alex to be. But notice that it's not really clear whether Billy wants to know where Alex is currently on sabbatical or where Alex teaches or where Alex went to graduate school or where Alex grew up. And—the point for us—Billy might not know what he wants to know. He just wants to know a bit more about Alex and will decide after she answers whether he got an answer to his question or not. He doesn't have to have the level of granularity sorted out before he asks the question. So context (or context plus Billy's intentions) need not resolve the contextual ambiguity.[8]

[5] "A Puzzle about Meaning and Communication," *op. cit.*, p. 350, renumbering the example.

[6] Buchanan, "A Puzzle about Meaning and Communication," *op. cit.*, p. 358.

[7] "'Might' Made Right," *op. cit.*

[8] "'Might' Made Right," *op. cit.*, p. 118, renumbering the example.

To handle such cases, they suggest,

> we can think of utterances taking place against a cloud of admissible contexts....
>
> There is no such thing as "the context," only the contexts admissible or compatible with the facts as they are.[9]

This cloud of admissible contexts will generate a cloud of propositions associated with the utterance.

Von Fintel and Gillies are interested in the cloudy picture because they think it can help with a puzzle about epistemic modals (like the 'might' in 'Joe might be in Boston'). These modals are standardly taken to be sensitive to a contextually supplied body of information: 'it might be that *P*' is true just in case

(10) $\text{Poss}_i\ P$
It is compatible with the information i that *P*,

where i is determined in context. However, recent critics of contextualist views have argued that in many cases, no one setting for this contextual parameter can explain both the readiness of speakers to make epistemic possibility assertions and the readiness of listeners to reject them.[10] For example, Alex can be warranted in asserting

(11) The keys might be in the car

even when, for all Alex knows, Billy may know that the keys are not there. To explain this, it looks as if we need to restrict the contextually

[9] Von Fintel and Gillies, "'Might' Made Right," *op. cit.*, p. 118.

[10] See Huw Price, "Does 'Probably' Modify Sense?" *Australasian Journal of Philosophy*, LXI, 4 (1983): 396–408 (considering probability rather than possibility); Andy Egan, John Hawthorne, and Brian Weatherson, "Epistemic Modals in Context," in *Contextualism in Philosophy*, ed. Gerhard Preyer and Georg Peter (Oxford: Oxford University Press, 2005), 131–68; Seth Yalcin, "Epistemic Modals," *Mind*, CXVI (2007): 983–1026; John MacFarlane, "Epistemic Modals Are Assessment Sensitive," in *Epistemic Modality*, ed. Andy Egan and Brian Weatherson (Oxford: Oxford University Press, 2011), 144–78; John MacFarlane, *Assessment Sensitivity: Relative Truth and Its Applications* (Oxford: Oxford University Press, 2014), chap. 10. The example (11) comes from von Fintel and Gillies, "'Might' Made Right," *op. cit.*, but it is similar to examples used by these relativist and expressivist critics of contextualism.

relevant information to what Alex knows (A), so that (11) expresses $\text{Poss}_A\, K$. But this yields the wrong predictions about the significance of agreeing or disagreeing with Alex's assertion. If Billy agrees with Alex, saying

(12) That's right, they might be,

she will normally be indicating, not that it is consistent with Alex's information that the keys are in the car, but that it is consistent with her own information (Billy's). This is shown by the fact that, if Billy knows that the keys are not in the car, it is inappropriate for her to agree with Alex:

(13) # That's right, but I know they're on the table.

Instead, she ought to disagree:

(14) No, they can't be in the car: I just checked there.

So, to agree or disagree with Alex's assertion of (11) is not to agree or disagree that it is consistent with what Alex knows that the keys are in the car. We can make sense of this by supposing that the information relevant to Alex's assertion of (11) includes what Billy knows. Perhaps it is Billy's information (B), or perhaps it is Alex's and Billy's pooled information (AB). But then it becomes hard to understand how Alex could have been warranted in asserting (11) in the first place. Alex was not in a position to know $\text{Poss}_B\, K$ or $\text{Poss}_{AB}\, K$. (And the difficulty increases if we consider eavesdroppers not known to the speaker.) There seems to be no one setting of the contextually relevant information that can explain both Alex's entitlement to assert (11) and Billy's entitlement to agree or disagree.

These observations have been taken to support relativist or expressivist alternatives to the standard contextualist view. Von Fintel and Gillies think this goes too far. They want to acknowledge the data while defending a contextualist semantics. To do this, they argue that there is indeterminacy about the context relevant to evaluating (11): the speech situation is compatible with a cloud of possible settings for

the relevant information state. Hence Alex's assertion cannot be identified with any *one* of the propositions one gets by plugging a particular body of information in for i in $\mathrm{Poss}_i K$. Instead, in uttering (11) Alex expresses a *cloud* of propositions, including $\mathrm{Poss}_A K$, $\mathrm{Poss}_B K$, and $\mathrm{Poss}_{AB} K$. Alex is warranted in making her claim because she has warrant for one of the propositions in the cloud ($\mathrm{Poss}_A K$). But Billy's response can target one of the other propositions ($\mathrm{Poss}_{AB} K$).

Though von Fintel and Gillies are mainly concerned with the case of epistemic modals, the way they motivate the cloud-of-propositions picture suggests that they take it to be an appropriate response to contextual underdetermination in general. The case they use to introduce the picture ("Where are you from?") has nothing to do with modals. And they concede that their strategy would seem ad hoc if the cloud-of-propositions view did not apply also in other cases, such as implicit quantifier domains.[11]

2.2 How to do things with clouds

So far we have been talking about *expressing* clouds of propositions. But in making a speech act, one does more than *express* contents. One asserts them, supposes them, asks about them. In order to understand the cloudy picture, then, we need to understand what it is to make an assertion (or other speech act) using a cloud of propositions. Here, as we shall see, the three views we are considering give very different answers.

Right away we face a terminological issue. The term 'proposition' is sometimes stipulated to mean the content of an assertion.[12] On this way of talking, the thesis that what is asserted is a cloud of propositions is nonsensical. It may turn out that what is asserted is a cloud of somethings, but if so, the cloud itself is the content of the assertion, and hence a proposition, and the somethings that compose it are something else.

[11] "'Might' Made Right," *op. cit.*, p. 123.

[12] Richard Cartwright, "Propositions," in *Analytical Philosophy*, ed. R. J. Butler, vol. 1 (Blackwell, 1962).

There are two ways around this terminological impasse. We can say that the speech acts one is making when one deploys a cloud of propositions are not assertions, or we can reject the idea that propositions are the contents of assertions.

Von Fintel and Gillies take the first approach. They talk of speakers "putting in play" clouds of propositions, in situations where they would not be entitled to "flat-out assert" all of them.[13] That looks like a denial that the speech act is one of assertion.[14] The terminology "putting in play" may be misleading. We might normally take ourselves to be "putting propositions in play" when we ask a question, make a conjecture, or perform many other kinds of illocutionary acts. But von Fintel and Gillies are using the term for a specific kind of speech act, whose force they will go on to describe. It might be better to introduce a new term, like 'cloudserting.' But as long as we keep firmly in mind that *putting in play* is supposed to be a specific, assertion-like speech act whose content is a cloud of propositions, we can avoid confusion.

Braun and Sider take the second approach: they say that assertions can have clouds of propositions as their contents. This commits them to rejecting the idea that propositions are the contents of assertions.

Buchanan is more difficult to pin down. He avoids the terminology of 'assertion' and 'speech act' entirely, talking instead of what the speaker means, so he does not take a stand on whether the speech acts in question are assertions. He says that propositions are the "objects of belief and certain other of our cognitive attitudes" and that they "determine truth-conditions."[15] It seems open to him to say either that some assertions have clouds of propositions as their contents, or to say that the speech acts in question are not, strictly speaking, assertions.

I do not think it matters much how these terminological issues are resolved. Whatever we decide to call the speech act in which we "put forward" clouds of propositions, we face the same substantive questions. How can we characterize the force of this speech act? What

[13] "'Might' Made Right," *op. cit.*, pp. 119–20.
[14] See also "'Might' Made Right," *op. cit.*, p. 117 n. 18.
[15] "A Puzzle about Meaning and Communication," *op. cit.*, p. 341.

norms govern it? How does it affect the common ground? How does it function in communication? What is required for "uptake?" The standard answers to these questions presuppose that the content of an assertion is a single proposition, so any proponent of the cloudy picture owes us a new story: a replacement for Classical Pragmatics.

2.3 Braun and Sider

One natural view is that in putting forward a cloud of propositions, one commits oneself to the truth of every proposition in the cloud. That is essentially the view of Braun and Sider.[16]

On Braun and Sider's view, vague sentences lack truth values, even when they do not involve borderline cases. For a sentence to be true in a context, they think, there must be a unique proposition that it expresses at that context, and that proposition must be true. The uniqueness condition fails for sentences containing vague expressions, even for non-borderline sentences like

(15) A person with no hair is bald.

So all such sentences lack truth values. Truth, then, "is an impossible standard that we never achieve."[17] However,

> ...it is usually harmless to *ignore* vagueness, set it aside, and act as if one's sentence is not vague, but rather expresses a unique proposition. When vagueness is being ignored, the cooperative communicator satisfies her communicative obligations well enough by uttering sentences that are approximately true...[18]

To say that a vague sentence is *approximately true* is to say that all of the propositions associated with it—the propositions that would be expressed by it on various legitimate disambiguations of its vague expressions—are true. For example, (15) is approximately true, because

[16] "Vague, so Untrue," *op. cit.*

[17] Braun and Sider, "Vague, so Untrue," *op. cit.*, p. 135.

[18] Braun and Sider, "Vague, so Untrue," *op. cit.*, p. 135.

on any legitimate way of disambiguating 'bald,' it will express a true proposition.

Thus, when we are ignoring vagueness (as we usually are in everyday life), the norm for asserting or "putting forward" a cloud of propositions is that each proposition in the cloud be true.[19] Asserting a cloud is committing oneself to the truth of every proposition in the cloud. And the effect of accepting such an assertion is the addition of all of these propositions to the common ground. This is a nice, simple story, which is compatible with Classical Contents and the standard conception of the common ground as a set of worlds. It is recognizable as a form of the popular supervaluational approach to vagueness.[20]

But this simple story is not plausible: the norm it proposes for vague assertion is too stringent. As Chris Barker[21] has noted, in a context where Richard's height is mutually known, a speaker might assert

(16) Richard is tall

with the aim of constraining the range of legitimate disambiguations of 'tall.' The force of the assertion is to propose that we go on in such a way that, within our conversation at least, Richard counts as tall. Such an assertion makes sense only if, prior to the assertion, there are legitimate ways of precisifying 'tall' on which Richard would not count as tall. But if that is so, the assertion is not permitted by Braun and Sider's norm of approximate truth. In allowing the assertion of (16) only in cases where it is already settled that Richard counts as tall, the norm appears to be too stringent.

[19] When we are not ignoring vagueness, on the other hand, *no* assertion of a cloud of propositions can meet the standard for assertion.

[20] Kit Fine, "Vagueness, Truth and Logic," *Synthese*, xxx, 3–4 (1975): 265–300; see Caie, "Semantic Indecision," *op. cit.* for an argument that this multi-proposition view is the only plausible way to interpret supervaluationism.

[21] "The Dynamics of Vagueness," *Linguistics and Philosophy*, xxv, 1 (2002): 1–36; Chris Barker, "Clarity and the Grammar of Skepticism," *Mind and Language*, xxiv, 3 (2009): 253–73; Chris Barker, "Negotiating Taste," *Inquiry*, lvi, 2–3 (2013): 240–57.

Though Braun and Sider do not consider this objection, it is one they might meet by appealing to *accommodation*.[22] Accommodation, as Lewis describes it, is the process by which nonfactual contextual parameters affecting the interpretation of utterance—which he conceives as "components of conversational score"—are adjusted when needed in order to interpret an utterance as "true, or otherwise acceptable."[23] For example, suppose the conversation is currently governed by a high "standard of precision," so that geometrical terms like 'hexagonal' apply only to figures with very straight sides, and someone asserts

(17) France is hexagonal.

The proposition (17) expresses given the current standard of precision is that France has six perfectly straight sides, and this is already ruled out by the common ground. So, in order to interpret the assertion as one that might be true, we will take the standard of precision governing the conversation to have been tacitly relaxed.[24]

In effect, Lewis's view posits a *bifurcated* common ground, consisting of a standard factual common ground—a set of worlds—plus a set of nonfactual "scoreboard" parameters that affect interpretation. When accommodation is triggered, the nonfactual components of the common ground are adjusted. When it is not, the factual component is adjusted. Accommodation is triggered in cases where the current settings of the nonfactual components would yield an update incompatible with the common ground.

By appealing to this mechanism, Braun and Sider could explain how an assertion of (16) can have the effect of contracting the range

[22] David Lewis, "Scorekeeping in a Language Game," *Journal of Philosophical Logic*, VIII, 3 (1979): 339–59.

[23] Lewis, "Scorekeeping in a Language Game," *op. cit.*, p. 347.

[24] Here is Lewis's "general scheme for rules of accommodation for conversational score": "If at time t something is said that requires component s_n of conversational score to have a value in the range r if what is said is to be true, or otherwise acceptable; and if s_n does not have a value in the range r just before t; and if such-and-such further conditions hold; then at t the score-component s_n takes some value in the range r" (Lewis, "Scorekeeping in a Language Game," *op. cit.*, p. 347).

of legitimate disambiguations of 'tall.' If there are legitimate disambiguations on which Richard does not count as tall, then (16) cannot meet the standard of approximate truth. So the participants in the conversation will naturally accommodate the speaker, adjusting the range of legitimate disambiguations on the conversational scoreboard so that Richard counts as tall on every legitimate disambiguation.

However, this reply only works in the case where Richard's height is mutually known. For only in that case do we know how to adjust the range of legitimate disambiguations of 'tall' so that Richard counts as 'tall' on all of them. Let us imagine, then, that Richard is not in the room for us to see. We have all seen him before, but we are unsure about his exact height: as far as the group knows, it could be anywhere between 190 and 195 cm. Suppose the prior range of legitimate disambiguations for 'tall' allows thresholds from 185 to 195 cm. What does our story say, now, about the force of asserting (16)?

In this scenario, the current range of legitimate disambiguations does not preclude the assertion's being approximately true. For if Richard is 195 cm—which is an open possibility given the common ground—then he exceeds the threshold for 'tall' on any legitimate way of disambiguating 'tall.' Since no change in the range of legitimate disambiguations is required for (16) to be approximately true, accommodation is not called for in this case. The update proposed by the assertion of (16) is that Richard is 195 cm tall—or so the framework predicts.

But this prediction is wrong. One can assert (16) without ruling out the possibility that Richard is a bit shorter that 195 cm. We might later discover, for example, that Richard is 193 cm tall, and accept an assertion to this effect. According to the theory we are considering, the common ground already excludes this possibility, so the context "crashes" and must be repaired. In addition, the earlier assertion of (16) must be regarded as false and retracted, just as an earlier assertion of

(18) Richard is 195 cm tall

would have to be. But in fact, we can learn that Richard is 193 cm tall without retracting (16), and without a crash in the context. Once we learn Richard's height, the earlier assertion of (16) commits us to changing the range of legitimate disambiguations for 'tall' to the range 185 cm to 193 cm, so that Richard counts as tall no matter how 'tall' is disambiguated. (16) is, then, a kind of conditional commitment: a resolution to contract the range of disambiguations for 'tall' as needed, when new information about Richard's height comes in. It is not, as Braun and Sider's theory would predict, an unconditional commitment about Richard's height.

One might object that accommodation can be triggered even when it is not strictly required for the *truth* of an utterance, if the utterance would otherwise be unacceptable in some other way. Lewis[25] gives an example of this kind: in a conversation with his friends in New Zealand, an assertion of 'the cat has gone upstairs' can trigger a reassignment of the *salient cat* from the New Zealand cat to the Princeton cat, not because the sentence could not be true of the New Zealand cat, but because the speaker in Princeton could not be *warranted* in making this claim about the New Zealand cat. One might argue that the same considerations force a revision of the set of legitimate disambiguations when (16) is asserted, if it is common ground that the speaker would not be warranted in asserting (18).[26]

But the problem is not, fundamentally, one of warrant. Even in cases where the speaker would be presumed to be warranted in asserting (18), it is not plausible that an assertion of (16) amounts to a commitment to (18), or that it would need to be withdrawn were we to discover (perhaps against the evidence) that Richard is not 195 cm tall. The commitment expressed by an assertion of (16) is best understood

[25] "Scorekeeping in a Language Game," *op. cit.*, p. 349.

[26] For example, it might be suggested that the top of the range of permissible thresholds for 'tall' is reduced to the largest d such that the speaker would be warranted in asserting that Richard's height is at least d. Is it plausible, though, that speakers and hearers would be able to converge on a value for d? That would require extremely fine-grained agreement on facts about the speaker's warrant, as well as agreement on how much warrant is required for the speaker to count as 'warranted.'

as a commitment to a certain relationship between the range of legitimate thresholds for 'tall' and the range of possible heights for Richard. The real problem is that there is no way to represent this kind of update in a bifurcated common ground, since it does not require any change in either the nonfactual scoreboard (the range of thresholds) or the factual scoreboard (the range of possible heights).

The point can be made even more simply by considering a disjunctive assertion:

(19) Either everyone taller than 193 cm is tall, or Richard is not taller than 193 cm.

As before, we assume that the nonfactual component of the common ground allows thresholds for 'tall' between 185 and 195 cm, and the factual component takes Richard's height to be between 190 and 195 cm. What update to the common ground is required if (19) is accepted? Not an update to the nonfactual scoreboard, because the second disjunct of (19) is compatible with any threshold for 'tall.' And not an update to the factual common ground, because the first disjunct of (19) does not rule out any possible height for Richard (or anyone else). So it seems that (19) does not require any update to the bifurcated common ground. Yet (19) is not a trivial assertion. One might have grounds for rejecting it, and accepting it constrains what one says in the future. We need a conception of common ground that allows (19) to make a nontrivial update.

Even if Braun and Sider's sort of supervaluationism worked for vague predicates, it would not be useful for the other cases of felicitous underspecification considered by Buchanan, King, and von Fintel and Gillies. Consider a variant of King's example (3). We look down the beach and say, of a group of surfers with indeterminate boundaries,

(20) Those guys are the *only* ones on the beach who really know how to surf.

Suppose again that 'those guys' has two legitimate interpretations:

(a) *those guys* = Abe, Bob, Cindy, Maria, Zeke, or

(b) *those guys* = Abe, Bob, Cindy, Maria, Sid.

If we were committing ourselves to (20) being true on every legitimate disambiguation of 'those guys,' then the commitment would be inconsistent. For (20) to be true on interpretation (a), it must be the case that Sid does not really know how to surf. But for it to be true on interpretation (b), it must be the case that Sid does really know how to surf.

Or consider von Fintel and Gillies' case of Alex, Billy, and the keys. If assertoric commitment is commitment to the truth of *all* the propositions in the cloud, then in asserting that the keys might be in the car, Alex is committing herself to the proposition that Billy's information does not rule out the key's being in the car. But that seems too strong. Alex is not in a position to make a claim about what Billy's information leaves open; she is asserting (11) partly to flush out any relevant information Billy might have.

We need a different story, then, about the force of asserting or otherwise "putting in play" a cloud of propositions.

2.4 Buchanan

Buchanan rejects the (Braun/Sider) view that when a speaker asserts a cloud of propositions (or, as he prefers to say, a *proposition type*), the intended update is to add the conjunction of the propositions to the common ground. The conjunction, he notes, "is simply another candidate proposition that Chet did not mean."[27] This is shown, he thinks, by the fact that Chet's audience "need not entertain each of [the propositions in the cloud] in order to understand the utterance."[28] It is sufficient for uptake that the audience "entertain *any one*, or more, of the candidates on the basis of the utterance," and be "thereby *disposed* to accept some number of the other salient candidates."[29] Thus,

[27] "A Puzzle about Meaning and Communication," *op. cit.*, p. 353.
[28] "A Puzzle about Meaning and Communication," *op. cit.*, p. 353.
[29] "A Puzzle about Meaning and Communication," *op. cit.*, p. 366, n. 22.

> *understanding a speaker's utterance requires entertaining some one or more propositions which are of the restricted proposition-type the speaker meant.* In the case of the utterance of ['Every beer is in the bucket'], Tim need not entertain the restricted proposition-type that Chet means; rather, what is required is that Tim entertain one or more of the candidates of that type on the basis of Chet's utterance.[30]

To understand Chet's assertion of (4), then, Tim need not grasp the proposition type (or cloud) Chet intended; he need only fasten on one of the propositions in this cloud. It does not matter which.

Buchanan is not specific here about how he sees the proposed update to the common ground, but what he says makes it clear that grasping the proposed update cannot require recognizing what proposition type (or cloud of propositions) the speaker intended. One hearer might fasten on one proposition, another on another. This is, for Buchanan, a *feature* of the view: he wants to explain how "an utterance might be understood in non-equivalent, yet equally correct ways."[31]

But this feature makes it impossible to see how the speech act could be viewed as proposing any specific update to a *common* ground. Suppose Chet asserts (4), and Tim and Zeke both signal their assent. What has been added to the common ground? It may be that Tim has fastened on the proposition that every beer they bought in the bodega is in the bucket, while Zeke has fastened on the proposition that every beer they will serve at the party is in the bucket. These are different updates to the common ground.[32] If there is no common knowledge as to which of these updates has been accepted, then we cannot think of them as adding to the common ground.

Buchanan presents himself as rejecting what might seem a peripheral part of Grice's account of speaker meaning—the idea that what is communicated is a proposition—while keeping the basic shape of the account. But his view commits him to rejecting the *central* idea

[30] *Ibid.*, p. 359.
[31] *Ibid.*
[32] Unless it is already common ground that they will only serve the beers they bought at the bodega.

of Grice's account: the idea that the speaker's meaning intention is *inter alia* an intention that the audience recognize this intention.[33] It is this essential transparency, Grice thinks, that distinguishes meaning intentions from other kinds of intentions to produce effects on hearers. And, as Robert Stalnaker[34] observes, it is the transparency of Gricean meaning intentions that allows us to think of speech acts as updating a common ground. Buchanan rejects the transparency of meaning intentions. On his view, the speaker's intention is satisfied if the audience entertains *some* proposition of the relevant type, but the speaker need not intend that the audience recognize this intention—for that would require recognizing the proposition type the speaker intended, and Buchanan denies that this is necessary for successful uptake.

Would it help to give up this part of Buchanan's view, and say that uptake requires recognizing the proposition type intended by the speaker?[35] A natural thought is that the speaker intends the hearer to recognize the type intended, but allows the hearer to pick *which* token of that type is to be added to the common ground. This would allow us to continue to model the common ground as a set of accepted propositions—but only if it becomes mutually known which proposition the hearer has fastened on to. This would happen if Tim responds to Chet's assertion of (4) by saying,

(21) Yes, every beer we bought at the bodega is in the bucket full of ice on the back porch,

thus clarifying the update. But what about the more normal case where Tim simply responds with 'Yep' or a nod? Then there will be no com-

[33] Buchanan seems to recognize that he needs to modify this part of Grice's view (see "A Puzzle about Meaning and Communication," *op. cit.*, p. 368 n. 34).

[34] "Common Ground," *Linguistics and Philosophy*, xxv, 5–6 (2002): 701–21, at p. 704.

[35] Perhaps Buchanan resists this because many of the same worries about determinacy that arose for propositions can be raised for proposition types. That is, there are many different proposition types—and accordingly many distinct clouds of propositions—that a speaker might be taken to have asserted. But let us leave this issue aside for now, since our task is to see what sense might be made of asserting a cloud of propositions.

mon knowledge about what the update is supposed to be. And that is just to say that we cannot view the update as an update to a *common* ground. For an account like this to succeed, we would need something like an algorithm for determining which proposition the hearer was responding to.

2.5 von Fintel and Gillies

The view sketched by von Fintel and Gillies is designed to solve this problem, giving clear predictions about how the common ground should be updated after various kinds of response to an epistemic possibility claim. Unfortunately, as I will argue, although it relies on principles that should be generally applicable if they are valid at all, it only gives good results in the specific case of epistemic possibility claims.

Recall that on von Fintel and Gillies' view, a bare epistemic modal claim like

(22) It might be that *P*

"puts in play" a cloud of propositions of the form

(23) It is compatible with information *i* that *P*,

with values for *i* taken from a contextually determined range. We have been asking what it is to "put in play" a cloud of propositions of this kind. What norms govern this speech act, and how does it affect the common ground?

In answer to this question, von Fintel and Gillies articulate two norms: one governing the making of this speech act and another governing its uptake (acceptance or rejection). The norm for making the speech act is

> Assert
> One may put in play a cloud of propositions just in case one is in a position to "flat out assert" one of the propositions in the cloud.[36]

[36] Von Fintel and Gillies, "'Might' Made Right," *op. cit.*, p. 120.

Two exegetical notes, before we continue. First, although von Fintel and Gillies state Assert as a narrow principle governing utterances of sentences of the form 'It might be that ϕ,' I have formulated it as a more general pragmatic principle. Our interest is in exploring what sense, in general, can be made of "putting in play a cloud of propositions." If Assert turns out to be plausible only for epistemic possibility claims, then one would want to look for a more generally applicable principle from which it follows, given special features of epistemic possibility modals. Otherwise Assert looks ad hoc.

Second, von Fintel and Gillies' wording suggests only a necessary condition:

> Suppose an utterance of *might(B)(ϕ)* by S puts in play the propositions $P_1, P_2, \ldots$. Then S must have been in a position to flat out assert one of the P_i's.

> Our proposal is that in order for a speaker to be within her linguistic and epistemic rights when she issues a BEM [bare epistemic modal sentence] against a cloud of contexts, she has to be in a position to flat out assert one of the meanings it can have, given that cloud.[37]

However, their argument requires not just a necessary but a sufficient condition for the permissibility of an assertion. Applying Assert to their scenario, they say: "given the facts of that scenario, Alex is justified in uttering the BEM *iff* she is justified in claiming that her evidence does not rule out the prejacent."[38] Their argument moves from Alex's entitlement to flat-out assert one of the propositions in the cloud, $\text{Poss}_A K$, to her entitlement to "put in play" the cloud. This move requires a *sufficient* condition for the permissibility of putting in play the cloud, not just a necessary condition. Accordingly, I have presented Assert as a necessary and sufficient condition.

Assert gives us one piece of the puzzle: it shows why Alex can be entitled to assert (11), even though she does not know what Billy knows about the location of the keys. What about the other piece? Why is

[37] *Ibid.*
[38] *Ibid.*, emphasis added.

Billy entitled to reject Alex's claim, on the basis of her knowledge that the keys are not in the car?

To sort this out, we need to understand what the hearer is supposed to do when a speaker "puts in play" a cloud of propositions. The basic idea is that the hearer selects one of the propositions in the cloud to react to, accepting or rejecting it. Thus, by putting in play a cloud of propositions, instead of a single one, the speaker cedes some control to the hearer in determining what update to the common ground is to be made. If the hearer selects the proposition P, and accepts it, then P is added to the common ground. If she selects P and rejects it, then the negation of P is added to the common ground.

For this story to work, though, it must be possible for all parties to the conversation to figure out which proposition has been targeted by the hearer. Otherwise, there will be no shared understanding about how the common ground is to be updated when the hearer accepts or rejects the speaker's claim. So, the hearer cannot be given complete freedom to target any proposition from the cloud:

> ...not just any one of them will do. Instead, we argue that the hearer is guided by what response to which proposition will be most informative in the conversation. When the modal is an existential like might, this will in fact lead to a dominance of negative replies.[39]

Von Fintel and Gillies articulate this constraint on the hearer's choices through a principle they call

> Confirm/Deny
> One may confirm (deny) a speech act that puts in play a cloud of propositions just in case one takes the strongest proposition in the cloud that one reasonably has an opinion about to be true (false).[40]

Let us see how this applies to the case of Alex and Billy. In asserting (11), Alex puts in play a cloud containing three propositions: $\text{Poss}_A K$, $\text{Poss}_B K$, and $\text{Poss}_{AB} K$. $\text{Poss}_{AB} K$ is logically stronger than $\text{Poss}_A K$:

[39] *Ibid.*, p. 121.

[40] *Ibid.* As with Assert, I have restated this as a general principle about cloudy speech acts, rather than bare epistemic modal claims.

if Alex and Billy's combined information is compatible with the keys being in the car, then Alex's information is too. So Confirm/Deny tells Billy to target $\text{Poss}_{AB}\, K$, provided she has a reasonably grounded opinion about its truth. Since Alex and Billy mutually know, in this case, that Billy is in a position to have an opinion about the truth of $\text{Poss}_{AB}\, K$, it becomes mutually known that this is the proposition Billy was targeting, and the common ground can be updated accordingly.

There is a puzzle here about how Confirm/Deny is supposed to relate to the norm of making the most informative response, to which von Fintel and Gillies appeal in motivating it. After all, accepting a stronger proposition is more informative than accepting a weaker one, but *rejecting* the stronger proposition is less informative than rejecting the weaker one. Yet Confirm/Deny says that the hearer should target the strongest proposition she has a reasonable opinion about, whether she is confirming or denying it. According to von Fintel and Gillies, the most informative move Billy can make is to reject $\text{Poss}_{AB}\, K$. But presumably Billy is also in a position to reject $\text{Poss}_{B}\, K$, which is weaker than $\text{Poss}_{AB}\, K$, and rejecting a weaker proposition would be more informative than rejecting a stronger one. If we stick to the "most informative move" motivation, then, we should predict that if Billy *accepts* Alex's claim, he is targeting $\text{Poss}_{AB}\, K$, while if he *rejects* it, he is targeting $\text{Poss}_{B}\, K$.

If we are going to take into account all of the possible responses Billy might make, we also need to be able to say whether rejecting the stronger proposition $\text{Poss}_{AB}\, K$ would be more informative than *accepting* the weaker proposition $\text{Poss}_{A}\, K$. In the scenario as described, accepting $\text{Poss}_{A}\, K$ is not very informative: it is probably already common ground after Alex's utterance that she does not know where the keys are. So, in this particular case, rejecting $\text{Poss}_{AB}\, K$ is more informative than accepting $\text{Poss}_{A}\, K$. But that does not follow from the logical relation between $\text{Poss}_{AB}\, K$ and $\text{Poss}_{A}\, K$; it depends on special features of the case. There are certainly cases in which accepting a logically weaker proposition would be more informative than rejecting a logically stronger one. For example, if the question at hand is who is

coming to the party, then accepting the proposition that Sarah is coming is more informative than rejecting the (stronger) proposition that Sarah and the entire population of Albuquerque are coming.

The upshot is that the instruction given in Confirm/Deny—"always target the strongest proposition in the cloud about whose truth value you have a reasonable opinion"—conflicts with the intuitive motivation of maximizing the informativeness of one's contribution. I am not sure whether to give precedence to the ideal of maximizing informativeness or the explicit statement of Confirm/Deny, so in what follows I will consider both options.

What is attractive about von Fintel and Gillies' proposal is that it offers a way of determining which proposition the hearer is fastening on, and hence how the common ground should be updated. Thus, if Billy rejects Alex's claim, the common ground can be updated with $\neg\,\mathrm{Poss}_{AB}\,K$.[41] On the other hand, if Billy accepts Alex's claim, then it becomes common ground that $\mathrm{Poss}_{AB}\,K$ is true. In either case, the exchange leads to a standard update of the common ground, conceived in the conventional way as a set of worlds: "once a hearer has confirmed or denied the BEM with all its indeterminacy, the resulting common ground is quite determinate."[42]

It is Confirm/Deny (or perhaps the ideal of maximizing informativeness) that allows the parties to reason in this way. Without this instruction for determining which proposition the hearer is targeting, von Fintel and Gillies' story would be much like Buchanan's, and it would suffer from the same flaw: it would be unclear how the common ground is to be updated after a claim is accepted or rejected.

The problem is that this view does not generalize well beyond the case to which von Fintel and Gillies apply it: bare epistemic possibility claims. The picture requires that the hearer can figure out in a transparent way which of the propositions in the cloud the hearer is targeting. Otherwise we would not have a well-defined update to the common ground. In the case of a bare epistemic 'might' statement, von

[41] *Ibid.*, pp. 123–24.
[42] *Ibid.*, p. 124.

Fintel and Gillies claim, we can assume that the hearer is targeting the strongest proposition in the cloud that she "reasonably has an opinion about." In our toy example, it is clear to both parties that Billy is in a position to have a reasonable opinion about $\text{Poss}_{AB}\, K$, and that this is the strongest proposition in the cloud. So Alex can come to know which proposition Billy is rejecting, and it is transparent how the common ground is to be updated.

But now consider a different sentence:

(24) The keys are probably in the car.

Our cloud now consists of the propositions $\text{Prob}_{A}\, K$, $\text{Prob}_{B}\, K$, and $\text{Prob}_{AB}\, K$, where '$\text{Prob}_{i}\, P$' means 'It is probable given the information i that P.' Suppose that $\text{Prob}_{A}\, K$ is true and known by Alex. Then, according to Assert, Alex is warranted in asserting (24). Now what about Billy? Should she accept or reject Alex's statement? According to Confirm/Deny, she should identify the strongest proposition in the cloud whose truth value she has a reasonable opinion about, and target that one. But in this case, none of our three propositions entails any of the others, so none is strongest.[43] If she *does* agree with or reject Alex's claim, then, it will not be clear how to update the common ground.

Thus it looks as if Confirm/Deny can only do the job it is meant to do—showing us how the common ground is to be updated after a speech act that "puts in play" a cloud of propositions is accepted or rejected—in the core case of epistemic possibility modals, and it only works there because of a special property of these modals that is not even shared by epistemic modals in general.

In this particular case, one might try to salvage things by (a) arguing that Billy is not in a position to have a reasonable opinion about

[43] This is because Poss is monotonic in a way that Prob is not. Adding information can only remove possibilities, not add them, so if K is possible given AB, it must be possible given A. By contrast, adding information might, depending on the case, make something either more probable or less probable. (Simple example: learning that a horse has won all its previous races will increase our subjective probability that it will win this race; learning that it has lost all its previous races will decrease it.) So we have no entailment either way between $\text{Prob}_{A}\, K$ and $\text{Prob}_{AB}\, K$.

$\text{Prob}_{AB}\, K$, and (b) forgetting about Confirm/Deny and reasoning informally about informativeness, instead of focusing on logical strength. Suppose it is common knowledge that the only propositions in the cloud Billy is in a position to have a warranted opinion about are $\text{Prob}_{A}\, K$ and $\text{Prob}_{B}\, K$. Since Alex already knows $\text{Prob}_{A}\, K$ but does not know the truth value of $\text{Prob}_{B}\, K$, it is more informative for Billy to target $\text{Prob}_{B}\, K$. So, as long as Alex knows all of this, and knows that Billy does not have an opinion on $\text{Prob}_{AB}\, K$, she can work out how the common ground is to be updated if Billy accepts or rejects her claim.

However, even if Billy is not in a position to have a reasonable opinion about $\text{Prob}_{AB}\, K$, it is hard to see how this could be common ground between Alex and Billy. After all, for all Alex knows, it may be that Billy's information is strictly stronger than Alex's. That is, for all Alex knows, it may be that Billy knows everything that Alex knows that is relevant to the location of the keys, plus more in addition, and it may be that Billy knows this fact. In that case, something would be probable given *AB* just in case it is probable given *B*, so Billy *would* be in a position to have a reasonable belief about $\text{Prob}_{AB}\, K$.

Nor do we have to go far to find other cases with multiple possible updates, none of which is strictly more informative than any of the others. Consider our variant (20) of King's case of the surfers on the beach. In this case there are two possible interpretations,

(a) Abe, Bob, Cindy, Maria, and Zeke are the only ones on the beach who really know how to surf, and
(b) Abe, Bob, Cindy, Maria, and Sid are the only ones on the beach who really know how to surf.

Neither of these interpretations entails the other. Nor is there any looser sense in which one would be a more informative contribution to the conversation than the other.

One might try to salvage the core of the cloudy view while throwing out the idea that there is an algorithm (such as Confirm/Deny) for determining which proposition the hearer is targeting. Instead, one might say that the hearer has an obligation, in accepting or rejecting

the speaker's claim, to make it clear to the speaker which proposition she is targeting. She might do this directly:

(25) No, it isn't probable, given what we both know, that the keys are in the car.

Or she might do it indirectly—for example, by backing up her claim with evidence that would only be relevant to one of the propositions in the cloud:

(26) No, it isn't likely to be in the car; I am pretty sure that I would have noticed it when I was in the garage just now.

In this way we could preserve the idea that the update is in some way a matter for negotiation between the speaker and hearer, while giving up the idea that both parties can figure out which proposition the hearer is targeting without explicit hints.

But I do not think this approach will help, in general, with the cases of felicitous underspecification that have typically motivated appeals to clouds. For when hearers accept or reject, they are typically *not* targeting a single, determinate proposition from the cloud, any more than speakers are.

Consider Buchanan's case (4). It is plausible that the speaker's intentions are compatible with a cloud of propositions that differ in exactly how the incomplete definite descriptions 'the beer' and 'the bucket' are supplemented. But the same goes for the *hearer*'s intentions. Chet asserts (4); Tim assents with a laconic 'Yep.' If we now ask Tim whether he was assenting to

(27) Every beer we bought at the bodega is in the bucket full of ice on the back porch, or
(28) Every beer we will serve at the party is in the bucket on the back porch, or
(29) Every beer in the apartment is in the bucket with pictures of pirates on it,

will he be able to answer? Even if he does answer, does his answer reflect a determinate intention he had when he assented to (4), or has he just plumped for something retrospectively? In this case, the indeterminacy left open by the speaker does not seem to be resolved in the hearer's response.

Or consider King's case of the surfers down the beach. Suppose King says

(30) Those guys are good.

and his hearers reject his claim, saying,

(31) No, they're not that good, they just got lucky on that wave.

Do his hearers need to have any more determinate an idea than King of who, exactly, belongs to the plurality denoted by 'they?'

The problem is perhaps clearest in the case of vague gradable adjectives. Suppose Anna sees Tim, who stands 1900 mm tall, and says,

(32) He's tall,

thereby putting in play a cloud of propositions of the form

(33) T_n = Tim is at least n mm tall,

for every n between 1850 and 2000. There is no plausibility to the idea that, when her hearers accept or reject Anna's claim, they are singling out a specific proposition in the cloud. Their intentions may be just as indeterminate as Anna's. Nor does Confirm/Deny (or the informal guidance to maximize informativeness that motivates it) help single out a specific proposition, even though the propositions in the cloud stand in clear relations of logical strength. Suppose Anna's hearers accept her claim. We are supposed to update the common ground with the strongest proposition in the cloud they reasonably believe to be true. Presumably it is reasonable for them to believe that Tim is at least 1850 mm tall, since that is already common ground. What about 1851 mm? 1852 mm? Even if there is a fact of the matter about what is the greatest n such that Anna's hearers reasonably believe that Tim is at least n mm

tall, it is not a fact we are in a position to know. So there is no way for the participants in this conversation to coordinate on a single proposition in the cloud with which to update the common ground.

In the general case, then, we cannot expect that a hearer's acceptance or rejection of a speaker's cloudy assertion will lead to convergence on an update to a standard, factual common ground (a set of worlds). That this can happen in the case of epistemic 'might' is due to special features of that case, which do not carry over even to other epistemic modals like 'probably.'

2.6 Toward a better solution

The cloudy picture is an attempt to understand how communication can work in the presence of underdetermination. Von Fintel and Gillies present it as a conservative alternative to the heretical relativist and expressivist views. It is conservative in preserving Classical Contents. But the price of this conservatism in the theory of content is radicalism in pragmatics: Classical Pragmatics must be given up and replaced by something new. We have looked at three different proposals for what this new picture could look like and found all of them lacking.

It is therefore worth considering an alternative way of modifying the Simple View of Communication: keeping Classical Pragmatics and abandoning Classical Contents. A first step is to allow propositions to vary in truth not just with a possible state of the world, but with one or more nonfactual parameters. In the case of vague gradable adjectives, this might be a delineation function that provides thresholds; in the case of epistemic modals, it might be an information state. If we do this, we can keep the idea that the content of an assertion is a single proposition, and that the proposed update is to add this proposition to the common ground. This immediately solves the major problem facing Buchanan's and von Fintel and Gillies' views: the lack of convergence on an update. For we now have a single content that is asserted by the speaker and recognized by the hearer.

We also solve the problem we saw with Braun and Sider's supervaluationist version of the cloud view. Recall that on this view, the standard for assertion is "approximate truth," or truth on every legitimate disambiguation. That led to the prediction that the upshot of accepting (16), in a context where it is common ground that Richard is between 190 and 195 cm in height and the range of legitimate thresholds for 'tall' is 185 to 195 cm, is to add to the common ground that Richard is 195 cm tall. Intuitively, though, to accept an assertion of (16) is not to rule out any particular height for Richard, but to commit ourselves to Richard's height being above the threshold for 'tall.' To make room for this kind of update, we need to think of the common ground as a constraint on *combinations* of worldly states of affairs and delineations. We can represent such a constraint as a set of (world, delineation) pairs—the combinations that meet the constraint. Once we do that, it is natural to model the content of (16) the same way. (Update, then, becomes simple set intersection.)

There remains a problem of interpreting this formalism. We know what it is to accept an orthodox proposition—one that is true or false at a possible world—but what is it to accept one of these souped-up contents? And what is it for a set of (world, delineation) pairs to be common ground in a conversation? Until these questions are answered, we do not know how to apply our formalism.

Chris Barker, who should be credited with the basic formal insight that updates constrain combinations of delineations and worldly states of affairs, spoils the insight by failing to treat the delineation as a nonfactual parameter. He describes ruling out delineations as reducing our "ignorance" and "uncertainty" about features of our discourse.[44] In fact, he reasons, since the discourse is a feature of the world, we do not actually need world/delineation pairs: we can just operate with worlds and a function, d, that returns the delineation of a given world.[45]

[44] Barker, "The Dynamics of Vagueness," *op. cit.*, pp. 3–4, p. 9.

[45] *Ibid.*, pp. 5–6; for further discussion, see Barker, "Clarity and the Grammar of Skepticism," *op. cit.*, p. 257; Barker, "Negotiating Taste," *op. cit.*, p. 246.

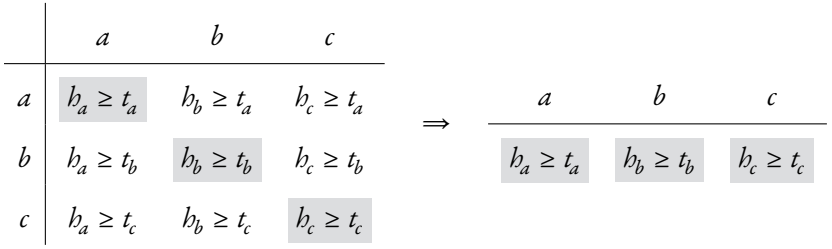

Figure 2.2: Diagonalizing interpretation of Barker's formalism. Here h_w = Richard's height in world w, and t_w = the threshold for 'tall' in world w. When we do not know whether we are in a, b, or c, diagonalization is triggered and we get the proposition that we are in a world w where Richard's height in w exceeds the threshold for counting as 'tall' in w. Note, however, that this interpretation is only possible if the participants presuppose that every situation determines a threshold for 'tall'.

On this interpretation, Barker's view reduces to the combination of epistemicism and diagonalization that we considered at the end of Lecture I (see Fig. 2.2). Asserting that Richard is tall communicates a condition on worlds: that Richard's height at the world is greater than the threshold for 'tall' governing the world. If this worked, the Simple Model would be vindicated. Contents and the common ground could be modeled as sets of worlds.

However, as we saw in Lecture I, diagonalization is only possible if the conversational participants presuppose that the diagonal is well defined. To presuppose this, in the case at hand, is to presuppose that each possible world determines a threshold for 'tall.' Do we normally presuppose this? I think not. If we did, we would find it perfectly appropriate, when we know Richard's exact height but consider him a borderline case of a tall man, to express our ambivalence using the language of subjective uncertainty:

(34) Richard might be tall.
(35) Richard is probably tall.

And we should find it completely inappropriate to use the language of choice:

(36) We could count Richard as tall.
(37) We should probably count Richard as tall.

But it is just the reverse: (36) and (37) sound fine in this context, while (34) and (35) would normally suggest that there is uncertainty about Richard's height. This suggests to me that, when there are multiple delineations in play, we do not think that this is because we are ignorant of some fact of the matter. We think, rather, that we have not yet made up our minds about a practical matter—where to draw the line for 'tall.'

If this is right, then we must not think of the delineation component of Barker's pairs as determined by the world component. We should think of contents and the common ground, then, as a set of pairs, with one element of each pair representing the content of a maximally determinate belief, and the other representing the content of a maximally determinate plan or intention. The idea should feel familiar: it is, more or less, the conception of contents we find in Allan Gibbard's plan expressivism.[46] These contents represent mental states that combine doxastic and practical elements, and that have, in general, a dual direction of fit—partly world-to-mind, partly mind-to-world.

Giving flesh to this idea will be our task in Lecture III.

[46] Allan Gibbard, *Thinking How to Live* (Cambridge, MA: Harvard University Press, 2003).

Lecture III: Indeterminacy as Indecision

What happens when a speaker makes a vague assertion, like

(1) Some of the tall test tubes contain hydrofluoric acid,

and a hearer understands it? Following Stalnaker, we can think of the assertion as a proposal to modify the common ground in some way. Understanding it (*uptake*) consists in recognizing how the speaker proposes to modify the common ground. If the hearer accepts the assertion, the common ground is updated in the way proposed. But this is schematic: to fill out the story, we need to say more about what the common ground consists of, and how an assertion of (1) proposes to modify (or *update*) it. That has been our concern in the previous two lectures.

The classical way of filling out this schema is to suppose that an assertion has a single proposition as its content, and this proposition is understood as a truth condition (the sort of thing that might be represented, for some purposes anyway, as a set of possible worlds). The update associated with an assertion that p is to add p to the common ground, conceived as a set of propositions (or for some purposes as a set of possible worlds, the ones that are open possibilities given what is accepted). But if we try to understand (1) in this way, we run into trouble. For, as I argued in Lecture I, (1) is flexible enough to be used to assert many different truth conditions. In a typical case, the speaker's intentions do not single out just one of the candidates, and apprehension of a single candidate does not seem to be required for understanding. This is a special case of the more general phenomenon of felicitous underspecification. The problem does not go away if we allow propositions themselves to be vague. Nor can it be resolved by diagonalizing.

It is tempting to suppose that an assertion of (1) is associated with a *cloud* of propositions—all those consistent with the speaker's intentions. But what, then, counts as understanding an assertion of (1), and how should the common ground be updated if the assertion is accepted? In Lecture II, we considered some answers to these questions from the literature and found them unsatisfactory. On Ray

Buchanan's view,[1] uptake is achieved no matter which proposition in the cloud the hearer fastens on. But how is the common ground to be updated? Since the speaker will not, in general, know which proposition the hearer has targeted, we cannot make sense of an evolving common ground. Kai von Fintel and Anthony Gillies attempt to solve this problem by giving a norm that identifies which proposition in the cloud the hearer should be targeting.[2] They show how in cases of bare epistemic possibility statements, speaker and hearer can deploy their common knowledge of this norm to coordinate on a common ground update. However, this story does not generalize to other cases of felicitous underspecification (and thus looks ad hoc in the narrow case where it works).

If there is no way to single out a single proposition from the cloud, we might, with Braun and Sider,[3] take the speaker to be proposing that *all* the propositions in the cloud be added to the common ground. This approach yields a version of supervaluationism. But, as we saw in Lecture II, this sort of view cannot account for assertions whose intended effect on the common ground cannot be modeled as a change either to a factual common ground (a set of possible worlds) or to a nonfactual scoreboard—including, for example, a range of *delineations* that set thresholds for the gradable adjectives.[4] Someone who asserts

(2) Richard is tall

in a context where it is known only that Richard's height falls inside a certain range need not be ruling out either any particular possible worlds or any particular thresholds for 'tall.' The point might be, rather, to link the question of which thresholds govern 'tall' to the question of Richard's actual height, in a way that might constrain future updates. (If we later find out that Richard is 192 cm tall, then we can

[1] "A Puzzle about Meaning and Communication," *op. cit.*

[2] "'Might' Made Right," *op. cit.*

[3] "Vague, so Untrue," *op. cit.*

[4] Following Barker, "The Dynamics of Vagueness," *op. cit.* and David Lewis, "General Semantics," *Synthese*, XXII, 1–2 (1970): 18–67, I define a *delineation* as a function from gradable adjective meanings to thresholds.

rule out thresholds for 'tall' above this height; and conversely if we later rule out thresholds for 'tall' above 192 cm, we will be committed to Richard's being at least 192 cm tall.) This kind of commitment cannot be represented as any determinate update to a bifurcated common ground consisting of a set of worlds and a range of delineations. To represent it, we need to represent the common ground not as a pair of sets—a set of worlds and a set of delineations—but as a set of world/delineation pairs. The update proposed by (2) is to rule out all world/delineation pairs (w, d) at which Richard's height in w falls short of the threshold established for 'tall' by d.

This is precisely the representation we find in Chris Barker's pioneering work on the dynamics of vague discourse.[5] But how is this formal representation of the common ground to be understood? We know what it is for a set of worlds to represent the common ground in a group: a world belongs to the set just in case it is not ruled out by the group's common knowledge. But what is it for a set of world/delineation pairs to represent the common ground in a group?

Barker's approach is to think of the delineation as a special kind of fact—a fact about the "discourse."[6] This leads Barker to conceive of assertions as reducing our "ignorance" or "uncertainty" about "the" threshold governing the discourse.[7] The epistemicist's ideology has returned. Indeed, so understood, Barker's view is exactly the view we considered at the end of Lecture I, combining the epistemicist view that the context determines a threshold with diagonalization.

We rejected that approach on the grounds that speakers and hearers do not presuppose that there *is* a contextually determined threshold. This means that they are not in a position to coordinate on a diagonal proposition as the asserted proposition. But if we do not think of the delineation as a feature of the world, but as something *nonfactual*—something not determined even by a complete enumeration of the truths about a world—what does it mean for a set of world/

[5] Barker, "The Dynamics of Vagueness," *op. cit.*

[6] Barker, "The Dynamics of Vagueness," *op. cit.*, pp. 5–6. Formally, the delineation is determined as a function of the world of evaluation.

[7] "The Dynamics of Vagueness," *op. cit.*, pp. 3–4, 9.

delineation pairs to be the common ground of a conversation? That is the question we will attempt to answer in this lecture. I will argue that Allan Gibbard's *plan expressivism* gives us a way of thinking about the common ground that makes sense of Barker's formal model. I will then consider the upshots of the resulting expressivist view for some of the traditional questions about vagueness: bivalence, the sorites paradox, higher-order vagueness, and the nature of vague attitudes. I will conclude with some remarks about the utility of vague language: why there is often a point to speaking vaguely even when more precise talk is available.

3.1 From uncertainty to indecision

Epistemicists like Williamson[8] argue that we have inexact knowledge of the locations of thresholds for vague gradable adjectives, and that it is impossible to know where these thresholds lie because knowledge requires a margin for error. If knowledge is impossible, that might also explain why we do not have full beliefs about the threshold's location. But *partial* belief—the kind of credence associated with epistemic uncertainty, and standardly represented using probabilities—should be fine. Perhaps I am not in a position to know that the threshold for 'tall' is at 185 cm or 190 cm, but I ought to be able to have a credence—say, 0.75—that it lies between these points. Consider standard cases of inexact knowledge. If I have inexact knowledge (by perception) of the height of a pole, I can say things like

(3) Miguel is probably taller than that pole.
(4) Miguel might be taller than that pole.
(5) I can't tell whether Miguel is taller than that pole.

So, if the epistemicist is right that we have inexact knowledge of the location of the threshold for 'tall,' it should also make sense to say, when Miguel is a borderline case of a tall man, that

(6) ?Miguel is probably tall.

[8] *Vagueness*, *op. cit.*

(7) ?Miguel might be tall.
(8) ?I can't tell whether Miguel is tall.

It is certainly true that when Miguel is borderline tall, we will be *ambivalent* about classing him as tall.[9] But it seems to me that we would not express our ambivalence using sentences like (6)–(8). We would more naturally say

(9) We should probably count Miguel as tall.
(10) We could count Miguel as tall.
(11) I can't decide whether to count Miguel as tall.

The epistemicist's explanation of our ignorance of the location of the threshold does nothing to explain our unwillingness to use the language of subjective uncertainty here. And there is a clear contrast with ordinary cases of inexact knowledge. We would definitely not say

(12) ?We should probably count Miguel as taller than that pole.
(13) ?We could count Miguel as taller than that pole.
(14) ?I can't decide whether to count Miguel as taller than that pole.

(12)–(14) cannot be used to express our uncertainty about the relative heights of the pole and Miguel, because this is not a matter for our decision. (If we did say these things, we would be understood to be conveying that some vagueness in 'taller than'—for example, in whether height is measured with or without shoes—affects whether Miguel counts as taller than the pole.)

The lesson I draw from these contrasts can be stated simply. Our ambivalent attitudes in borderline cases should be understood as a kind of *practical indecision*, not as doxastic uncertainty. What we need, then, is a way of understanding the delineation component of the pairs in our integrated common ground as relating to *decision* in the way that the world component relates to belief.

[9] For this terminology, see Stephen Schiffer, *The Things We Mean* (Oxford: Oxford University Press, 2003), ch. 5.

3.2 Plan expressivism

The fundamental idea of Allan Gibbard's plan expressivism is that the states of mind we express using declarative normative sentences are not uniformly doxastic. They are, at least in part, *planning* states. Planning states relate to practical decision: they have a world-to-mind direction of fit, as contrasted with the mind-to-world direction of fit of standard doxastic states. The conclusion we came to in the previous section—that vague statements are partly doxastic, partly practical—suggests an approach to vague language along similar lines.

Our representation of the common ground as a set of world/delineation pairs is formally similar to Gibbard's representation of the contents of mental states as sets of world/hyperplan pairs.[10] *Hyperplans* are fully determinate contingency plans covering every possible circumstance and resolving all indecision. They are *hyper* because ordinary plans do not do this. Ordinary plans cover only a small range of circumstances (*I will pull over if I see a tornado from the highway*), and they leave a lot indeterminate (*pull over how fast? after activating the turn signal for how many seconds? in what sort of place?*). Still, we can represent an ordinary plan as the set of hyperplans that are compatible with it, just as we can represent the content of an ordinary belief as the set of possible worlds (which we can think of as fully opinionated belief states) that are compatible with it. An indeterminate and partial plan can be *firmed up*, by reducing the indeterminacy and extending the range of circumstances it covers. The firmed-up plan is compatible with the original plan, in much the same sense that the belief that Sammy is in London is compatible with the belief that Sammy is in England.

The reason Gibbard represents the contents of mental states as sets of world/hyperplan *pairs* is that many mental states are neither pure planning states nor pure doxastic states. The easiest way to see this is to suppose that D expresses a pure doxastic state and P a pure planning state, and then consider the state expressed by their disjunction,

[10] *Thinking How to Live*, *op. cit.*

$D \vee P$. The disjunction expresses a mixed state, which cannot be factored out into separate planning and doxastic components. To judge D is to rule out some possible worlds. To judge P is to rule out some hyperplans. But to judge $D \vee P$ is not to rule out any particular worlds or any particular hyperplans. Rather, it is to rule out a certain *combination* of further commitments: rejecting all the worlds compatible with D while rejecting all the hyperplans compatible with P.

But not all sentences expressing "mixed" states are disjunctive in form. Indeed, most ordinary normative judgments will not be pure planning states. Consider Gibbard's own example:

(15) I ought to pack.

To judge (15) is to take oneself to be in one of the situations for which one's plans call for packing. We can represent its content as

$$\{\langle w, h\rangle \mid h \text{ calls for packing in situation } w\}.$$

This is a constraint on world/hyperplan *pairs*; it does not rule out any particular worlds or any particular hyperplans.

My suggestion, then, is that we treat the delineation in our world/delineation pairs as a fully specific plan for setting a threshold for gradable adjectives.

This approach presupposes that thresholds for gradable adjectives are subject to planning. I take this to be the case. We can decide to regard only people above 192 cm in height as *tall*, or to treat only people who can pass a certain test as *qualified*. This is a matter for practical deliberation, decision, and planning. Normally, our plans are fairly indeterminate and do not single out a unique threshold for any gradable adjectives, let alone all of them, but only constrain the range of thresholds. In addition, the way they constrain thresholds is often linked to questions about the state of the world: for example, when we plan to count everyone at least as tall as Richard as tall, we link the practical question of where to set the threshold with the factual question of Richard's height. We can represent this constraint as the set of world/delineation

pairs such that Richard's height in the world exceeds the threshold for 'tall' determined by the delineation.[11]

When we think of the common ground as a set of world/delineation pairs, then, we are thinking of it as a complex joint commitment conditionally linking what is accepted about the world with joint plans for using words and concepts.

This conception suggests a natural generalization. I argued in Lecture I that the problem about vagueness and communication was a special case of the more general problem of felicitous underspecification. For gradable adjectives, we will need delineations in addition to worlds; for epistemic modals, we will need information states; for plural definites, we will need group boundaries; and so on. But we can think of all of these things as determined by a hyperplan. After all, a hyperplan is a maximally specific plan, one that encodes plans for everything—including, *a fortiori*, these aspects of the use of language. Though Barker was wrong to suggest that a possible world determines a delineation, a hyperplan does determine a delineation, and much else.

Having made this generalization, we can represent the common ground as a set of world/hyperplan pairs: a Gibbardian judgeable content. The problem of interpreting our formalism, then, reduces to the existing problem of interpreting the doxastic-practical states posited by plan expressivists.

On this view, a sentence like (2) has a fixed content (once the reference class[12] has been fixed):

(16) $\{(w, h) \mid$ the height of Richard in $w \geq$ the threshold for *tall* determined by $h\}$.

[11] Of course, different thresholds may be relevant for different uses of 'tall' in a discourse: someone who is tall for an academic may not be tall for an athlete. One way to handle this is to say that a delineation maps the semantic value of a gradable adjective to a threshold, but allow that 'tall' can have a different semantic value when used with different reference classes. In Kennedy, "Vagueness and Grammar," *op. cit.*, for example, the semantic values of gradable adjectives are degree functions, and phrases like 'for an athlete' modify these functions by restricting their domains.

[12] Or domain of the degree function.

The significance of an assertion of this content will depend on the state of the common ground:

- If it is already common ground that Richard is 189 cm tall, then an assertion of (16) will be tantamount to a proposal to plan to set the threshold for *tall* no higher than 189 cm.
- If it is already common ground that we plan to set the threshold for *tall* at exactly 189 cm, then an assertion of (16) will be tantamount to a proposal to add the factual proposition that Richard is 189 cm tall to the common ground.
- But if the common ground is agnostic about Richard's height (taking it to be between 185 and 195 cm) and undecided about the threshold for *tall* (not excluding options between 180 and 200 cm), then the update proposed by an assertion of (16) will be neither a plan nor a factual proposition. Rather, it will be a conditional commitment that ties together plans and factual beliefs.

This view agrees with standard "contextualist" and "interest-relative" views of vagueness[13] that thresholds (or threshold ranges) governing vague words can shift as a conversation unfolds. It disagrees with these views in taking this shift to be a change in the common ground rather than the contents asserted. Even if the threshold has shifted, you can assert the same content by (2) now that you would have earlier.

3.3 Bivalence, excluded middle, and indifference

Epistemicism is often presented as the view of vagueness one must have if one wants to retain bivalence and classical logic. And indeed, the main alternatives to epistemicism reject these: multivalued theories and fuzzy logics reject both bivalence and classical logic; supervaluationism retains classical logic but rejects bivalence. Where does our

[13] For example, Fara, "Shifting Sands," *op. cit.*; Scott Soames, *Understanding Truth* (Oxford: Oxford University Press, 1999).

expressivist view stand on these traditional questions? That depends on how we formulate bivalence, and on how we think of our hyperplans.

One way to formulate bivalence is

> Bivalence 1
> If S can be used to make a literal assertion at a context c, then either S is true at c or S is false at c (that is, the negation of S is true at c).

Expressivism rejects Bivalence 1, because according to expressivism sentences do not have truth values relative to a context. A context does not in general determine a unique delineation, and thus does not fix extensions for vague predicates or truth values for vague sentences. So Bivalence 1 is untrue. That is not to say that we accept "truth value gaps." Rather, the expressivist rejects the ideology presupposed by Bivalence 1: the classification of sentences as true or false at a context.

But there are other formulations of bivalence that are compatible with expressivist ideology. Following Williamson,[14] we can define monadic truth and falsity predicates, applicable to utterances, as follows:

(17) If an utterance u says that ϕ, then u is true iff ϕ and false iff $\neg\phi$.

As Williamson acknowledges, the predicates 'true' and 'false' in (17) are vague: if Sam is a borderline case of a tall man, then an utterance of 'Sam is a tall man' will be a borderline case of a truth.[15] Using these predicates, we can formulate bivalence as the schema

> Bivalence 2
> If u is an utterance that says that ϕ, then u is either true or false.[16]

Given (17), Bivalence 2 is equivalent to a restricted version of the principle of excluded middle:

[14] *Vagueness*, *op. cit.*, p. 188.

[15] Williamson, *Vagueness*, *op. cit.*, p. 192.

[16] The 'says that ϕ' part is needed, because no plausible version of bivalence would apply to utterances that cannot be described as 'saying that ϕ' for some declarative sentence ϕ (sighs, questions, meaningless noises, and so on).

(18) If u is an utterance that says that ϕ, then ϕ or $\neg\phi$.

Should an expressivist accept (18)? It seems to me that this issue hinges on how we think of the hyperplans. The way I have presented the expressivist view so far presupposes that a hyperplan determines a unique threshold for each gradable adjective.[17] This means that maximally decided states are never "indifferent" between several options for threshold-setting. On this assumption, the expressivist view validates the law of excluded middle and hence Bivalence 2. For, if *every* hyperplan in the common ground tells us to set a particular threshold for 'tall,' then the sentence

(19) Richard is tall or Richard is not tall

will be accepted in every common ground, because at every world/hyperplan pair, either Richard's height in the world exceeds the threshold for *tall* established by the hyperplan or it does not.

The fact that there is an expressivist view that accepts Bivalence 2 shows that there can be no good argument from Bivalence 2 to epistemicism. But it is easy to see why people have thought that there is one. Imagine a sequence of utterances

(20) At one second old I was young.
At two seconds old I was young.
At three seconds old I was young.
(And so on.)

If each of these utterances is either true or false, then surely there is a first false one, and hence a last second of my youth! How can we accept that without yielding everything to the epistemicist?

To dispel the puzzle, we need to see that on the expressivist view, 'true' and 'false' are plan-dependent. If one plans to use 'young' in a certain way, this plan will also affect one's use of the monadic predicates 'true' and 'false.' So, a common ground that accepts that an utterance

[17] More properly, each gradable adjective meaning (allowing for contextual sensitivity to a reference class): see n. 11, above.

u is either true or false may fail to accept either that *u* is true or that it is false: *u* may count as true on some possible ways of firming up our plans and false on others. What the expressivist concedes is that on any way of maximally firming up plans for using 'young,' there will be a last second of my youth. But that is not to say that, given the present indeterminate state of our plans, there is a fact of the matter about which second that is. It is this last point that the epistemicist needs us to accept, but Bivalence 2 is compatible with rejecting it.

So far we have argued that the expressivist can accept both excluded middle and Bivalence 2, on the assumption that each hyperplan puts the threshold for *tall* in a particular place. But what about that assumption? We can be undecided about where to draw the line between the tall and the non-tall, just as we can be undecided about whether to buy Crest or Colgate. But could we not also *decide* that it really does not matter where we draw the line between the tall and the non-tall—just as we might decide to be indifferent between Crest and Colgate? An *indifferent* state of mind would differ from an undecided one in disagreeing with a state of mind that plumps for one threshold (or one brand of toothpaste).

One might think that indifference in plans does not make sense: a plan is a plan *to do something* in a certain circumstance, and it has to select a particular course of action. (Of course, that course of action could be "flip a coin"—a course of action that is incompatible with choosing Crest or choosing Colgate, not indifferent between them.) But Gibbard's own expressivism allows for fully decided states that are indifferent between options. A Gibbardian hyperplan maps each circumstance, not to an option, but to a nonempty *set* of permissible options, between which it is indifferent. This allows Gibbard to distinguish the *undecided* state of mind

(21) being agnostic about which toothpaste one ought to buy, Crest or Colgate

from the *decided* but *indifferent* state of mind

(22) taking it to be permissible to buy either Crest or Colgate (that is, holding that it is not the case that one ought to buy Crest or that one ought to buy Colgate).

The content of (21) contains hyperplans that map the present circumstance to {buy Crest} and hyperplans that map it to {buy Colgate}, while the content of (22) contains only hyperplans that map the present circumstance to {buy Crest, buy Colgate}. Allowing hyperplans to be indifferent between options is crucial for Gibbard's aim of giving an expressivist account of normative vocabulary, since it allows him to solve the "negation problem"[18] and distinguish between (21) and (22). But it is also potentially problematic for his aim of giving a naturalistic account of normative states of mind. For what does it mean to say that a hyperplan is indifferent between several courses of action? Intuitively, it means that the plan regards all of them as permissible. But to say that is to describe the plan's contents in normative terms. If we have to think of plans as judgments that certain acts are permissible, then Gibbard has not really succeeded in explaining normative judgments in terms of more basic psychological states.[19]

We are not trying to give an expressivist account of normative judgment, and the negation problem does not arise for gradable adjectives the way it does for modals.[20] So nothing about the use to which we want to put plan expressivism compels us to countenance hyperplans

[18]Gibbard, *Thinking How to Live*, *op. cit.*, pp. 54–56; James Dreier, "Relativism (and Expressivism) and the Problem of Disagreement," *Philosophical Perspectives*, XXIII, 1 (2009): 79–110; Mark Schroeder, *Being for: Evaluating the Semantic Program of Expressivism* (Oxford: Oxford University Press, 2008).

[19]I am indebted here to an insightful unpublished paper, "Preference and the New Negation Problem," by Sophia Dandelet.

[20]The negation problem arises because we need to distinguish between three states of mind we can describe by putting negations in different places: (i) not believing one ought to ϕ, (ii) believing it is not the case that one ought to ϕ, and (iii) believing one ought not to ϕ. Without indifferent hyperplans, the expressivist only has resources to describe two of these: (i) not planning to ϕ and (iii) planning not to ϕ. But with gradable adjectives, we do not have so many places to put negations. We only need to distinguish between (i) not believing that a six-foot-tall man is tall and (ii) believing that a six-foot-tall man is not tall. (i) corresponds to not planning to count a six-foot-tall man as tall, and (ii) corresponds to planning to count a six-foot-tall man as tall.

that are indifferent between several thresholds. For our purposes, we could think of hyperplans as fully determinate conditional intentions, which map circumstances onto particular actions.

On the other hand, nothing prevents us from allowing hyperplans to be indifferent between multiple delineations. If we go that way, there might be reason to reject Bivalence 2. It would be natural to say that, relative to a hyperplan that is indifferent between thresholds for 'tall' that are below Richard's height and thresholds that are above Richard's height, the sentence (2) is neither true nor false but has a third, intermediate truth status. One way of dealing with this third status would be to use a three-valued logic (such as Strong Kleene) that rejects excluded middle. Contents will then be represented as functions from world/hyperplan pairs to multivalues $(0, 1, \frac{1}{2})$. For example, the content of (2) would be

$$
\begin{aligned}
(23)\ f(w,h) &= 1 \quad \text{if } \forall t \in T_h\ r_w \geq t \\
&= 0 \quad \text{if } \forall t \in T_h\ r_w < t \\
&= \tfrac{1}{2} \qquad \text{otherwise} \\
\text{where } T_h &= \text{ the thresholds for } \textit{tall} \text{ permitted by } h \\
r_w &= \text{ Richard's height in } w.
\end{aligned}
$$

An expressivist who develops the view this way would reject excluded middle and hence Bivalence 2. However, this rejection of classical logic is not motivated by the desire to make sense of borderline cases or solve the sorites paradox. The failures of excluded middle and bivalence stem from indifference, which, if it exists, is distinct from indecision. As I will argue shortly, the latter is enough by itself to make sense of the phenomena distinctive of vagueness.

3.4 The sorites paradox

We have been ignoring the main preoccupation of the classical literature on vagueness—the sorites paradox—in order to get clearer about the role of vague language in communication. Let us now see what our view has to say about arguments like this:

(24) If people whose net worth is n dollars are rich, so are people whose net worth is $n - 1$ dollars.
(25) People whose net worth is a billion dollars are rich.
(26) So people whose net worth is one dollar are rich.

The expressivist view rejects the conclusion of this argument by rejecting its first premise. For expressivists who do not reject classical logic, this means accepting the negation of (24):

(27) There is an n such that people whose net worth is n dollars are rich but people whose net worth is $n - 1$ dollars are not rich.

So far this looks just like what the supervaluationist says about the sorites. But accepting (27) is less problematic for the expressivist than for the supervaluationist. A supervaluationist takes (27) to be *true*, despite the fact that there is no particular value n that yields a true instance. And this is hard to swallow. How can an existential claim be true if none of its instances are? An expressivist has no comparable problem. To say that (27) is accepted in a common ground is just to say that all the hyperplans we are undecided between put the threshold for 'rich' *somewhere*. And there is no difficulty understanding how an existential claim can be *accepted* in a common ground even when none of its instances are. For example, it can be common ground that somebody in the class ate the cookies, even though it is not common ground that Bill did, or that Sally did, or...

An account of the sorites argument must do more, though, than explain why the argument is unsound. It must also explain why it *seems* so compelling. Why is (24) so hard to reject, if its negation is accepted in the common ground?

The attractiveness of (24) stems from our sense that vague adjectives like *large* are *tolerant*, in Crispin Wright's sense—insensitive to small changes in a thing's size.[21] We cannot imagine classifying a person as rich and then refusing to classify a person who has just one dollar less as rich. And we take this little thought experiment to support (24).

[21] "Language Mastery and the Sorites Paradox," *op. cit.*

As Delia Graff Fara, Scott Soames, and others have pointed out, this is bad reasoning.[22] For there may be an explanation for our inability to conceive of a one-dollar boundary between the rich and the non-rich that is consistent with the existence of such a boundary. According to Fara and Soames, calling someone rich shifts the extension of 'rich' so that people with similar net worths fall into the extension of 'rich,' too. This explains why we cannot think of a counterexample to (24), without giving us any reason to suppose that (24) is true.

I agree that the appeal of (24) rests on a false inference from the difficulty of thinking of a counterexample. But, unlike Fara and Soames, I do not think this is because it is incoherent to accept that someone who has n dollars is rich while refusing to accept that someone with $n - 1$ dollars is rich.[23] On the account I have been sketching, it is possible for conversational participants to adopt a joint plan to impose a sharp one-dollar threshold for 'rich.' Nothing about the meaning of 'rich' prevents this,[24] and we can imagine cases where it might be useful to adopt such a plan: for example, if we have been told to separate out the rich applicants, we are required to classify every applicant either as rich or non-rich, and we are given the necessary information about their net worths. I gave my first talk on this topic in Cardiff, in

[22] Fara, "Shifting Sands," *op. cit.*; Soames, *Understanding Truth*, *op. cit.*

[23] Soames' idea that the boundary of the extension shifts by some unknown delta when we call something rich is objectionable for the same reason as epistemicism: it makes it impossible to have common knowledge of the update that is proposed. In discussing Soames' view, Timothy Williamson, "Soames on Vagueness," *Philosophy and Phenomenological Research*, LXV, 2 (2002): 422–28 rightly asks what we have gained by moving from one hidden boundary to another.

[24] This claim is controversial. Jerry Fodor and Ernie Lepore, "What Cannot Be Evaluated Cannot Be Evaluated and It Cannot Be Supervalued Either," this JOURNAL, XCIII, 10 (1996): 516–35 argued that when words like 'tall' are used with stipulated sharp thresholds, they are technical terms with special meanings. And Chris Kennedy argues that "the positive form [of a gradable adjective] cannot be used to distinguish between objects that are very similar relative to some gradable property" ("Vagueness and Grammar," *op. cit.*, sec. 2.3); he takes this to be a semantic constraint. I am arguing that the "crisp judgment" data that motivates Kennedy—the oddity of calling a book 'the long book' when it is only one page longer than another—can be explained pragmatically on the expressivist view, so we need not build constraints into the semantics that explain them.

the shadow of a hill called The Garth. The Garth is the subject of a film which recounts how villagers carried material to its top in order to push it over the 1,000-foot threshold for counting as a mountain in official maps of Britain.[25] Promontories on maps must be labeled as mountains or hills, and for these purposes it is necessary to agree on some cutoff point, however arbitrary.

If nothing prevents drawing a sharp line between the rich and the non-rich, how can we explain our intuitive repugnance to instances of (27)? On the view we have been exploring, accepting an instance of (27) is tantamount to accepting a fully determinate plan for using the concept *rich*, one that leaves no room for further refinement in light of experience or needs. In most contexts, we will have good practical reason to reject such plans. And certainly, in the artificial context of most sorites arguments, we lack compelling reasons to accept such proposals.

Similar considerations can explain our willingness to be drawn along in "forced march" sorites sequences. Suppose we have accepted that someone with n dollars counts as rich. The claim that someone is rich is, in part, a proposal about how to use 'rich' (or the concept *rich*). In accepting that someone with a net worth of at least n dollars is rich, one has accepted such a proposal. We are now asked to consider a new proposal: to count people with a net worth of at least $n - 1$ dollars as rich. In imagining that we have accepted the first proposal, we take for granted that it was a reasonable one, given our purposes. Could we nonetheless have some reason to reject the second proposal? Some kind of special reason would be needed, and one can imagine cases where there is one. But discussions of the sorites take place in a contextual vacuum. Nothing is said about our interests or purposes in classifying things as rich. So we search in vain for a reason to resist the proposal, and, unable to find one, we are drawn to accept it.[26]

[25] *The Englishman Who Went up a Hill but Came down a Mountain* (Miramax Films, 1995); Christopher Monger, *The Englishman Who Went up a Hill but Came down a Mountain: A Novel* (Miramax Books, 1995).

[26] See MacFarlane, "Vagueness as Indecision," *op. cit.*, sec. VII.

The basic point here is a general one about the rationality of planning, and it can be illustrated using examples that do not involve the application of vague gradable adjectives. Suppose we are planning a kid's birthday party. You say, "Let's get at least 15 balloons." I agree. A bit later, you say, "Actually, let's get at least 16." You have made a practical proposal, one that firms up our plans a little bit. I could reject your new proposal, but it is hard to think of reasons to reject it, given that I have accepted a plan to get at least 15 balloons.[27] Of course, if I am faced with a sequence of such proposals, I will eventually have to reject one—or drown in balloons. But the psychology here is much the same as in a "forced march" sorites from red to orange patches, where participants eventually get off the bus feeling that they have misclassified many orange patches as red.[28] And if I am right that vague assertions are partly proposal-like, the two cases can be explained in exactly the same way.[29]

3.5 Higher-order vagueness

Williamson's case for epistemicism rests largely on a general "revenge" argument: approaches that seek to avoid a sharp boundary between the tall and the non-tall inevitably seem to reintroduce sharp boundaries elsewhere.[30] For example, supervaluationists avoid commitment to an inscrutable threshold by considering a range of legitimate thresholds, but the endpoints of the range would seem to be just as inscrutable as the original threshold. If we are going to be committed to hidden semantic boundaries anyway, why not simply accept them at the first stage?

[27] It is true that if we imagine specific sorts of situations—perhaps balloons come in packs of five, or perhaps the room is so small that 15 balloons will completely fill it—then I may have good reason to balk. But the same is true of the sorites paradox, in light of the sort of cases (mentioned above) where it makes sense to agree on a sharp boundary.

[28] See Diana Raffman, *Unruly Words: A Study of Vague Language* (Oxford: Oxford University Press, 2014), ch. 5.

[29] See MacFarlane, "Vagueness as Indecision," *op. cit.*, sec. VII.

[30] *Vagueness*, *op. cit.*

On the expressivist view, as we have seen, the contents asserted do not depend on the range of legitimate thresholds.[31] So, if indeterminacy in the boundaries of the thresholds poses a problem, it is not in identifying the update that is being proposed to the common ground. It is, rather, a problem in identifying the common ground itself. A proposition p is in the common ground just in case it is common knowledge among the conversational participants that p is accepted.[32] If someone asserts

(28) Richard is tall,

we know that we are supposed to remove from the common ground all world/delineation pairs such that Richard's height in the world does not meet the threshold for 'tall' established by the delineation. That is something speaker and hearer can easily coordinate on. But we will only end up with common knowledge of the *result* of the update if we had common knowledge of the *prior* state of the context. Before (28) was asserted, which world/delineation pairs were still in play, and which were excluded?

In certain cases this question may have a clear answer. As a conversation unfolds, constraints can accumulate in a way that pins down the live possibilities. Suppose, for example, that it is common knowledge prior to the assertion of (28) that the shortest person who had been called 'tall' in the conversation was 195 cm, and that the tallest person who had been called 'not tall' was 180 cm, and that Richard is between 185 and 192 cm in height. Then it is clear that the common ground allows all combinations of heights h for Richard and thresholds t for 'tall' such that $185 \leq h \leq 192$ and $180 < t \leq 195$. But what about a case in which (28) is asserted near the beginning of a conversation, in which nobody else has been called 'tall?' What was the largest threshold for 'tall' not excluded by the prior common ground? Presumably a threshold of 160 cm was excluded. What about 161 cm? 162 cm? As

[31] See p. 69, above.

[32] Or true common belief, in the version of Stalnaker, "Common Ground," *op. cit.*, p. 716. The difference between common belief and common knowledge does not matter for the arguments I make here.

we move up, we will be progressively less confident that the value was excluded by the prior context.

But is there really a problem here? Suppose that 182 cm is the largest height such that it is common ground that anyone this height or taller counts as 'tall.' If there is a common ground at all, there must be *some* number that fits this description. Of course, it is not plausible to suppose that it is *known* that 182 cm is the number that fits this description. But why should it be known? All that is required by the notion of common ground is that there *be* a greatest height x that is ruled out as a threshold for 'tall' by the common ground. It is not required that it be common ground that x is the greatest such height.

By the definition of common knowledge, we can infer that if it is common ground that anyone of height $\geq x$ counts as tall, it is common ground that it is common ground that anyone of height $\geq x$ counts as tall. This is an instance of

> Positive Introspection
> If it is common knowledge that p, it is common knowledge that it is common knowledge that p.

Thus, if x is the bottom of the range of allowed thresholds for 'tall' in the context, it must be common ground that x is *in* the range. But it need not be common ground that x is the *bottom* of the range. For that, it would have to be common ground that it is *not* common ground that anyone of height $x - \partial$ counts as tall. But common knowledge does not satisfy

> Negative Introspection
> If it is not common knowledge that p, it is common knowledge that it is not common knowledge that p.[33]

It would be fallacious, then, to argue as follows: if there were a common ground of the sort posited, it would have to be commonly known where the range of allowed thresholds begins—but this is not commonly known, so there is no common ground.

[33] See *ibid.*, p. 707.

Indeed, if this were a good form of argument, it would strike just as much against a purely factual notion of common ground. What is common ground at the beginning of a conversation? It is presumably common ground, for example, that there have been black dogs, and that $2 + 2 = 4$. But is it common ground that Nixon resigned the presidency? One can start asking hard questions about where, exactly, common ground leaves off. However, an ability to answer these questions does not show that there is not common ground. Because it is not generally common ground that something *is not* common ground, the boundaries of the common ground are not in general going to be common knowledge.

The failure of Negative Introspection allows us to accept that we do not always know where the range of allowed thresholds begins and ends, while still holding onto the idea that there is a determinate such range.

3.6 Vague thought

The idea that vagueness is, at root, a kind of practical indecision may seem less plausible when we turn our focus from communication to thought. One might grant that when we communicate with each other, we have joint plans for the use of vague words such as 'tall,' and that these plans both shape and are shaped by the evolving conversation. But can this really be the essence of vagueness? After all, in addition to *saying* that Richard is tall, one can *think* this, even when one is alone. One might suppose that there is no role, in these cases, for plans for the use of words. Yet the thought that Richard is tall is vague. So vagueness cannot, at heart, be a matter of indecision. Focusing on the use of vague language simply distracts from the core phenomenon.

I want to resist this objection. It seems to me that there is a planning element even in thought. To recognize this, it is not necessary to say that we think in words. It is enough if our classificatory and inferential thought involves the application of *concepts*. In thinking about Richard, I can deploy a concept *tall* whose boundaries I recognize as

subject to decision and planning. Knowing his height in centimeters, and the heights of others in relevant classes, I can decide to classify Richard as tall or not in my own thinking. How I make this decision will depend on the other uses to which I am putting the concept *tall* in this episode of my thought. If I have decided to choose only tall people for my team, then deciding how to constrain the threshold for counting as tall is, in part, deciding what kind of team I want. We may not always be explicitly aware of the element of planning and decision that goes into thought, but we can easily bring it to attention.

When we engage in joint thinking with others, where inferences are to be drawn from premises supplied by different thinkers, we will need to *coordinate* our planning for the application of concepts. In practice, this is usually done via language. I give voice to my thought that Richard is tall by saying 'Richard is tall.' When everyone synchronizes their *tall* concepts with their use of 'tall,' it becomes possible to engage in joint planning and deliberation. It is not, then, that thoughts intrinsically involve plans for the use of words; rather, our plans for the application of concepts come to be linked with plans for the use of words when we need to coordinate with others.

One might still object: what about the thoughts of primitive creatures? Must we hold that thoughts can be vague only if the creatures who have them are capable of planning and deliberation about their own conceptual activity? Fish and frogs engage in a certain kind of classification, which many philosophers and cognitive scientists have wanted to regard as conceptual representation. Often the contents ascribed to such primitive thoughts are contents we would express using vague sentences: for example, *that is food*, or *there is danger*. Are these attributions unwarranted if we are not willing to ascribe to fish and frogs the ability to deliberate about where to draw the line between food and non-food, or between danger and non-danger?

I think it is a mistake to suppose that the classifications that are useful in thinking about human thought are always going to be applicable to more primitive forms of thought. I am willing to bite the bullet and say that it does not make sense to say that the frog's or the fish's

thoughts are vague. But I do not think that means we must say they are precise. Rather, the distinction between vagueness and precision is one we can sensibly draw only for creatures who can evince the sort of evidence brought out by sentences like (3)–(14). We can use vague sentences to ascribe thoughts to primitive creatures, but we need to keep in mind the limitations of doing this.[34]

On the view I have been defending, what is distinctive of vague thought is the element of practical decision or planning it involves. Vague thought is not fully doxastic; it has an element that is deliberative, practical, and "meta-contextual."

This account has the resources to answer an important question raised by Stephen Schiffer:[35] what, exactly, is it to *take something to be borderline*? When we consider Miguel as borderline tall, we are ambivalent about whether he is tall. And this ambivalence comes in degrees: we would be less ambivalent if Miguel were closer in height to a clear case of a tall or a non-tall man. It is tempting to assimilate this ambivalence to ordinary doxastic uncertainty, and that is what the epistemicist view requires. Schiffer argues that this is a mistake. Suppose we are ambivalent about whether to classify Miguel as tall, as bald, as smart, and as funny, because he is a borderline case on all four dimensions. If our ambivalence were a matter of doxastic uncertainty, then we should be pretty confident that Miguel is *not* tall, bald, smart, *and* funny—for in order to have this conjunctive property, Miguel must exceed the thresholds for all four individual properties. Assuming, as seems plausible, that the positions of the thresholds are probabilistically independent, the probability that Miguel exceeds all four should be *much* smaller than the probability that he exceeds any of them singly. We should be pretty confident, then, that Miguel is *not* tall, bald, smart, *and* funny. But in fact, Schiffer observes, we are about as ambivalent about the conjunction as we are about the conjuncts. He concludes

[34] A bolder way of making this point is that epistemicism might be just fine as an account of the frog's *food* thoughts. That does not show that it is fine as an account of our *food* thoughts.

[35] Stephen Schiffer, "Two Issues of Vagueness," *The Monist*, LXXXI, 2 (1998): 193–214; Schiffer, *The Things We Mean*, *op. cit.*, ch. 5.

that our ambivalence in borderline cases is not a kind of doxastic uncertainty.[36]

Well, then, what is it? It seems apt to call it a kind of partial belief: at any rate, someone who takes Miguel to be borderline bald does not fully believe either that he is bald or that he is not. Like subjective credence, this attitude admits of degrees, and Schiffer proposes that we represent the strength of this "vagueness-related partial belief" (VPB) using real numbers between 0 and 1, with 1 corresponding to full belief. But VPB differs from subjective credence or "standard partial belief" (SPB) in a number of important respects:

- When we have a VPB that p, it seems misguided to speculate about whether p.
- VPBs do not generate likelihood beliefs, like "there is at least a fifty-fifty chance that Joe is bald."
- VPBs can be held even when we think we are in the best possible epistemic position to pronounce on the truth of a proposition.
- VPBs do not seem to rationalize betting behavior like SPBs do.
- VPBs are not constrained by the probability calculus in the way that SPBs are, but rather by the Łukasiewicz rules for fuzzy logic (in which the degree of a conjunction is the minimum of the degrees of the conjunctions). This explains our ambivalence, noted above, toward the conjunctive proposition that Miguel is tall, bald, smart, and funny.

All of these contrasts are plausible, but they serve rather to pick out a subject matter than to explain it. *Why* does vagueness-related partial belief differ from uncertainty-related partial belief along these dimensions? Schiffer gives no answer: he offers no intrinsic characterization of VPB, just the list of contrasts with SPB. Nor does his theory

[36]This argument cuts not only against classical epistemicism, but against views like that of Bacon, *Vagueness and Thought*, *op. cit.*, who gives a non-epistemic account of vagueness but takes us to have standard credences over vague propositions. See John MacFarlane, "Andrew Bacon, *Vagueness and Thought*," *The Philosophical Review*, CXXIX, 1 (2020): 153–58.

have any coherent account of mixed states of mind, which involve both vagueness-related and uncertainty-related partial belief. For example:

> Mixed Case
> You do not know how many hairs Miguel has, but you have a normal (bell-shaped) distribution of credences over different possible numbers between 100 and 600. You regard men with 100–400 hairs as borderline cases of bald men, but you are more ambivalent toward classifying a man as bald if he has close to 400 hairs than if he has close to 100 hairs.

If asked whether Miguel is bald, you would be reluctant to answer, or you would hedge. But your reluctance is not just due to uncertainty about the number of hairs Miguel has, nor is it just due to indeterminacy about where the threshold lies. Your attitude toward the proposition that Miguel is bald is, rather, a mix of vagueness-related ambivalence and doxastic uncertainty. Schiffer's approach is to say that you have some VPB and some SPB that Miguel is bald, and some VPB and some SPB that Miguel is not bald. He insists that these four numbers—the positive and negative SPB and VPB—must sum to 1. But this constraint does not sit well with his claim that SPBs are governed by probabilistic constraints and VPBs by the Łukasiewicz rules, and in fact Schiffer's methods for analyzing mixed cases lead to contradictions.[37]

I want to suggest that the expressivist view gives us a more satisfactory account of vague thought than Schiffer's theory of VPBs, capturing most of what Schiffer was after, while also giving an elegant account of mixed cases that involve elements of both vagueness-related and uncertainty-related partial belief. On this view, one's attitude toward propositions one takes to be borderline is characterized by *practical indecision*. That is why it is misguided to speculate about the truth of such a proposition: its truth is still undetermined because we have failed to make a decision that would allow it to be determined, not be-

[37]John MacFarlane, "The Things We (Sorta Kinda) Believe," *Philosophy and Phenomenological Research*, LXXIII, 1 (2006): 218–24; John MacFarlane, "Fuzzy Epistemicism," in *Cuts and Clouds: Vagueness, Its Nature, and Its Logic*, ed. Richard Dietz and Sebastiano Moruzzi (Oxford: Oxford University Press, 2010), 438–63.

cause there is some fact out there of which we are ignorant. That is also why taking p to be borderline is compatible with thinking that one could not be in a better position to pronounce on the truth of p.

The reason vagueness-related ambivalence does not give rise to likelihood beliefs is that it does not make sense to assign probabilities to propositions like *Miguel is bald* until we have decided on a threshold for *bald*. If we fix a threshold, then the probabilities we assign to various worldly states of affairs—for example, to Miguel's having exactly 252 hairs on his head—will determine a probability for *Miguel is bald*. But while the threshold remains undecided, we have at best a range of likelihoods.[38] There is no way of collapsing this range to a single number.

One might think we can assign probabilities to our various possible threshold decisions, and multiply through. For example: if one thinks it is 80% likely that the conversation will evolve in a direction that sets a threshold for counting as *bald* below Miguel's number of hairs, one will consider it 80% likely that Miguel is bald. But it is a mistake to conflate the question how likely it is that one will plan to put a threshold in a particular place with the question how likely it is that the threshold is in that place. There is no making sense of the latter, as long as the issue is left undecided.[39]

[38] If probabilities are defined for sets of worlds, then we can derive a function from hyperplans to probabilities: where $P \subset W \times H$ is a Gibbardian content,

$$f(h) = Pr(\{w | \langle w, h\rangle \in P\}).$$

From this function we can derive a set of probabilities (the ones to which any hyperplans are mapped), or an upper and lower bound. But both moves lose information.

[39] We can illustrate the point by considering a similar move one might make with conditional preferences. Mary has two great loves—farming and the arts—and she is at a point where she must decide to which to devote her life. She prefers to live in the country if she devotes her life to farming, and prefers to live in the city if she devotes her life to the arts. Suppose she thinks she is 80% likely to decide to devote her life to the arts. Does that mean it is 80% likely that she prefers city living to country living? No. Since she has not yet decided whether to devote her life to farming or the arts, it is similarly undecided whether she prefers city living or country living. So there is no fact of the matter about which we might have a credence of 0.8. She has only conditional preferences about where to live, not unconditional ones.

For similar reasons, vagueness-related ambivalence does not rationalize betting behavior. If we are ambivalent toward the proposition that Miguel is tall because we have not decided where to put the threshold for counting as tall, it is senseless to accept a bet at even odds that Miguel is tall. We could bet on whether we will eventually set the threshold such that he counts as tall—but that is not the same as betting on whether he is tall.

One might worry that the expressivist view loses something that seemed right about Schiffer's view: the idea that vagueness-related partial belief admits of degrees. Schiffer dramatizes this by considering how one's degree of endorsement of the proposition that Tom Cruise is bald increases as hairs are plucked one by one from Cruise's head. But we need not represent VPBs as totally ordered, as Schiffer does, to recognize these differences. We recognize that every way of firming up our plans that would make a 150-haired man count as bald would make a 100-haired main count as bald, too, but not vice versa. That is, we could make a decision that classified 100-haired men as bald while leaving it open whether 150-haired men count as bald. This kind of inclusion relation yields a partial ordering for vagueness-related ambivalence and allows us to say that, in some sense, we have a stronger degree of commitment to the proposition that Tom-with-100-hairs is bald than the proposition that Tom-with-150-hairs is bald. These comparative considerations are enough to establish that one should not endorse a conjunction of vague propositions more strongly than the conjuncts. But they do not go further than that—which is why in Schiffer's example it is fine to endorse the conjunction to about the same degree as the conjuncts. We cannot extend these comparisons of degree beyond cases where we have this kind of inclusion relation (comparing one's attitude to *Tom is bald* to one's attitude to *Tom is smart*, for example), or move from the ordering to the kind of quantitative relationships we have with probabilities.

So far, I have been emphasizing that vagueness-related ambivalence is not to be confused with uncertainty-related credence. But surely, one might object, we *do* assign credences to vague propositions! For

example, I may think it is more than 50% likely that the lobby of my hotel will contain some soft chairs, or that the bodega on the corner will have fresh fruit. Our indecision about where to put the thresholds for counting as soft or as fresh does not pose much of an obstacle to these judgments. But that is because, in these cases, the effect of indecision about thresholds is small. Resolving our indecision about a threshold for *soft* one way rather than another will affect our estimation of the probability of soft chairs only very slightly. To put it another way, our indecision constrains the probability to a narrow range. If we like, we can think of this range as a single "imprecise probability."[40] In these cases, we can safely act as if the probability is somewhere in that range, as long as our actions do not depend on *where* it is in the range.

Indeed, in many cases, we can assign sharp probabilities to vague sentences, because indecision about the threshold only matters for states of affairs we have already ruled out. For example, if our plans constrain the threshold for counting as tall to between 185 and 195 cm, and have a 0.5 credence that Miguel is exactly 180 cm tall and a 0.5 credence that he is exactly 200 cm tall, then we should have a sharp 0.5 credence that he is tall.

In sum, then, what is distinctive of vague thought is its "plannishness"—the way in which it expresses practical decisions about threshold drawing. This feature explains why vagueness-related partial belief is distinct from standard uncertainty-related partial belief in the ways Schiffer describes. And it makes it easy to think about the "mixed cases" that are problematic for Schiffer. Indeed, Gibbard's formalism of sets of world/hyperplan pairs was designed precisely

[40] See, for example, Isaac Levi, "On Indeterminate Probabilities," this JOURNAL, LXXI, 13 (1974): 391–418; James M. Joyce, "A Defense of Imprecise Credences in Inference and Decision Making," *Philosophical Perspectives*, XXIV, 1 (2010): 281–323; Susanna Rinard, "A Decision Theory for Imprecise Probabilities," *Philosophers' Imprint*, XV, 7 (2015): 1–16. These authors are primarily concerned with cases where we have no epistemic basis for choosing between various possible probability distributions, rather than cases where a probability distribution depends on plans that we have not yet firmed up. However, the same formalism could be used for both purposes, and many of the same considerations apply.

to handle the similar issue of judgments with mixed normative and descriptive content.

3.7 Why be vague?

I want to close by discussing a question that is not, I think, asked enough. The question is this: why do we use vague words and concepts at all, when we have precise alternatives? Why should we describe Richard as 'tall,' for example, instead of saying that he is 'more than 180 cm tall?'

The problem is especially acute for the epistemicist. On the epistemicist's view, the word 'tall' (in context) expresses a property we could have expressed in precise terms—but we do not know *which* property. What reason could we have to use a term whose extension is unknown to us in preference to one whose extension is known to us?

The answer cannot be that vague terms are useful when we do not know the heights of things precisely. For even when we do not know something's precise height, we often do know that it falls within a precisely delimited range: for example, more than 180 cm tall, or between 180 and 200 cm. Given that we can ascribe these precise height ranges with full justification, why should we ever prefer the vague word 'tall,' which differs from them mainly in the obscurity of the lines it draws? How could our communicative or cognitive goals be furthered by using the more obscure term?

To be sure, 'tall' is shorter, but that is hardly an answer. If it were useful to do so, we could have introduced shorter synonyms for descriptions of precise ranges. And 'tall' is contextually sensitive, but that is not an answer either. Why do we not use a precise term that means, for example, "is at least 10 cm taller than the average of all items in the class *C*," where *C* is supplied by context?

I want to propose that the expressivist view developed here offers a compelling solution to the problem. In a nutshell: vague words are useful because they facilitate joint planning. Our joint endeavors of-

ten require us to draw lines and make distinctions. For example, in hiring people we need to distinguish the qualified from the unqualified; in designing filtering software we need to distinguish the pornographic from the non-pornographic. Getting agreement about where the lines should go can be difficult. But it is often quite easy for a group to agree about the consequences of drawing a line in a particular place, and on various constraints on where the line should go. A lot of planning can take place without settling where the line should go. Vague words make it easy for us to do this joint planning. We can begin to use a word like 'qualified' or 'pornographic' without deciding where the threshold is. As our deliberations advance, we can impose further constraints on the threshold ("this, at least, is not pornographic") and decide on consequences of exceeding it ("children will not be allowed to view pornographic images"). In most cases, our objectives can be met without coming to any final decision on a threshold, and even in the rare cases when such a decision is needed, it is helpful to be able to begin deliberations before making it. Vague words serve as tags or placeholders around which we can build consensus when it is difficult to draw precise lines.

Inquiry would be impoverished if we heeded Voltaire's advice: "Define your terms, I say, or we will never come to agreement."[41] Voltaire was right to warn us about the danger of merely verbal disagreements. But, especially when our goal is action rather than philosophical knowledge, requiring precisification too early can hinder coming to agreement. Better to leave lines undecided while we achieve consensus on the constraints governing them and the consequences of drawing them. Often, we will find, the consensus we are able to achieve either makes agreement on a precise threshold unnecessary, or makes it easier.

[41] "Définissez les termes, vous dis-je, ou jamais nous ne nous entendrons." *Dictionnaire Philosophique*, vol. IV, Oeuvres Complètes (Paris: Hachette, 1960), p. 205.