

Word and Plan

Columbia Themes in Philosophy

Columbia Themes in Philosophy

Series Editor: Akeel Bilgrami, Johnsonian Professor of Philosophy,
Columbia University

Columbia Themes in Philosophy is a new series with a broad and accommodating thematic reach as well as an ecumenical approach to the outdated disjunction between analytical and European philosophy. It is committed to an examination of key themes in new and startling ways and to the exploration of new topics in philosophy.

Edward Said, *Humanism and Democratic Criticism*

Michael Dummett, *Truth and the Past*

John Searle, *Freedom and Neurobiology: Reflections on Free Will, Language, and Political Power*

Daniel Herwitz and Michael Kelly, eds., *Action, Art, History: Engagements with Arthur C. Danto*

Michael Dummett, *The Nature and Future of Philosophy*

Jean Bricmont and Julie Franck, eds., *Chomsky Notebook*

Mario De Caro and David Macarthur, eds., *Naturalism and Normativity*

Alan Montefiore, *A Philosophical Retrospective: Facts, Values, and Jewish Identity*

Thom Brooks and Martha C. Nussbaum, eds., *Rawls's Political Liberalism*

Noam Chomsky, *What Kind of Creatures Are We?*

Word and Plan

John MacFarlane

Columbia University Press
New York



Columbia University Press

Publishers Since 1893

New York Chichester, West Sussex

cup.columbia.edu

Copyright © 2026 John MacFarlane

All rights reserved

Cataloging-in-Publication Data available from the Library of Congress

ISBN 9780231212809 (hardback)

ISBN 9780231212816 (trade paperback)

ISBN 9780231559577 (EPUB)

ISBN 9780231556483 (PDF)

LCCN 2025044878

Printed in the United States of America

Typeset in TeX Gyre Schola by the author.

Cover design: Tim Jones

Cover image: “*Impressio Librorum*,” *Nova Reperta* (Phillips Galle, c. 1590), plate 4. Engraving by Jan Collaert I, after Johannes Stradanus. From the collection of the Metropolitan Museum of Art. Public domain.

GPSR Authorized Representative: Easy Access System Europe,
Mustamäe tee 50, 10621 Tallinn, Estonia, gpsr.requests@easproject.com

In memory of my mother, Jackie MacFarlane

Contents

Preface	ix
1 Vagueness and Communication	1
1.1 Frege's View	5
1.2 Epistemicism	8
1.3 Excursus: This and That	12
1.4 The Context Sensitivity of Vague Language	14
1.5 Felicitous Underspecification	20
1.6 Why Not Diagonalize?	22
1.7 Conclusion	27
2 Seeing Through the Clouds	29
2.1 Motivating Clouds	31
2.2 How To Do Things With Clouds	36
2.3 Braun and Sider	37
2.4 Buchanan	43
2.5 von Fintel and Gillies	45
2.6 Toward a Better Solution	54
3 Indeterminacy as Indecision	59
3.1 From Uncertainty to Indecision	62
3.2 Plan Expressivism	64
3.3 Bivalence, Excluded Middle, and Indifference	67
3.4 The Sorites Paradox	72
3.5 Higher-Order Vagueness	76
3.6 Vague Thought	78
3.7 Why Be Vague?	85
Appendix: On King's Theory	87

4	Indecisive Semantics	95
4.1	The Task of Semantics	96
4.2	Degree Semantics for Gradable Adjectives	101
4.3	Semantics Without Truth at a Context	106
4.4	Have Meanings Gotten Too Thin?	109
4.5	Beyond Threshold Indeterminacy	111
4.6	Panvariabilism	116
4.7	Assignment-Sensitive Contents	121
4.8	Epistemic Shift	125
4.9	Semantic Shift	128
4.10	Adieu Diagonalization	131
4.11	Conclusion	133
5	Metalinguistic Negotiation	137
5.1	Plunkett and Sundell's Account	138
5.2	How Metalinguistic Disagreements Can Be Substantive	142
5.3	Missing Agreement	143
5.4	Inscrutable Contents	145
5.5	Confused Speakers	148
5.6	Mixed Cases	150
5.7	A Diagnosis	152
5.8	An Alternative	155
5.9	Half-Measures	160
5.10	Conclusion	170
6	Probabilistic Information	173
6.1	Information and Communication	175
6.2	The RSA Model	175
6.3	Applying RSA to Vague Language	181
6.4	A Conceptual Objection	185
6.5	Fixing the Bug	187
6.6	Further Problems for RSA Models	188
6.7	An Alternative	192
	Bibliography	199
	Index	209

Preface

Philosophers have traditionally been interested in vagueness as a source of paradox. Intuitively, removing a single grain of sand from a heap still leaves you with a heap—yet repeating this procedure will eventually leave you with a single grain. Must we then accept that a single grain can be a heap, or that removing a single grain can make a heap into a nonheap? Or can we avoid these repugnant conclusions by revising some principles of logic?

In this book, I will be concerned with a different, much less discussed problem: explaining how we can communicate using vague language. Communication is successful when a speaker says something and a hearer grasps what the speaker means. But when Abe says “Sam is tall” and Bertha understands, what exactly has been communicated? Some information about Sam’s height, it is natural to think. But what information? I argue here that standard approaches to vagueness have no good answer to this question. To understand communication with vague words, I argue, we need to see that what is communicated is not a piece of information in the normal sense. Vague talk is not purely factual; it is, in part, expressive of practical plans. The view is a form of the plan expressivism originally developed by Allan Gibbard as an account of normative language.

Chapters 1–3 were originally given as my Woodbridge Lectures at Columbia University and published in the *Journal of Philosophy*.¹ I am

¹ John MacFarlane, “Lecture I: Vagueness and Communication,” *Journal of Philosophy* 117, no. 11/12 (2020): 593–616; John MacFarlane, “Lecture II: Seeing Through the Clouds,” *Journal of Philosophy* 117, no. 11/12 (2020): 617–42; John MacFarlane, “Lecture III: Indeterminacy as Indecision,” *Journal of Philosophy* 117, no. 11/12 (2020): 643–67.

grateful to Columbia's Philosophy Department for inviting me to give the lectures, and to Akeel Bilgrami for inviting me to publish them in book form, together with some supplementary material. I have opted to make only minor changes to the original lectures. The only major addition is a discussion of Jeffrey King's *Felicitous Underspecification*, which was published after the lectures; to preserve the shape of the original lectures, this has been included as an appendix to chapter 3.

Three new chapters flesh out the picture by discussing issues I did not discuss in the lectures themselves, although many of them came up in the discussions following the lectures. Chapter 4 considers what the pragmatic picture argued for in chapters 1–3 implies for semantics. It defends the radical view that we should treat all lexical items as variables. Chapter 5 considers the relation of my proposal to David Plunkett and Timothy Sundell's influential work on "metalinguistic negotiation." It argues that the phenomena Plunkett and Sundell seek to understand are better explained by the expressivist theory developed here, and that metalinguistic negotiation, properly understood, does not threaten the traditional "argument from disagreement." Chapter 6 argues against a tempting alternative solution to the problem of vague communication, according to which vague assertions communicate probabilistic information.

My thinking on all these issues goes back to my 2016 talk at the Joint Session in Cardiff.² I spent the 2016–2017 academic year at the Institut d'Études Avancées (IEA) in Paris, with an office in a seventeenth-century mansion on the Seine, and I am very grateful to the staff and the other fellows for the warm intellectual environment they provided.³ I appreciated and profited from the challenge of formulating what I was up to in a way that could be appreciated by those outside of philosophy and semantic theory. In Paris, I also had valuable discussions with my colleagues at l'Institut Jean Nicod. My ideas began to take shape in my contribution to a workshop on contextual indeterminacy and semantic theory that I organized with François Recanati at the IEA in 2017 and were refined in

² John MacFarlane, "Vagueness as Indecision," *Aristotelian Society Supplementary Volume* 90, no. 1 (2016): 255–83.

³ This was made possible by the financial support of the French state managed by the Agence Nationale de la Recherche, program "Investissements d'avenir" (ANR-11-LABX-0027-01 Labex RFIEA+), and the Fondation Maison des Sciences de l'Homme.

two subsequent graduate seminars at Berkeley. chapter 6 had its origin in my contribution to a workshop on meaning and probability that I organized with Paul Égré in 2022, and chapter 5 as my contribution to the Kadish Workshop in Law, Philosophy, and Political Theory at Berkeley, hosted by Véronique Munoz-Dardé and Joshua Cohen. I am grateful to participants in these workshops and seminars, and to audiences at talks given in Chicago, Konstanz, Paris, Berkeley, Syracuse, Princeton, Stanford, Chapel Hill, and San Diego, for their helpful feedback. I was especially glad, on two of these occasions, to discuss my ideas with two Chrises whose important earlier work influenced mine: Chris Barker and Chris Kennedy. It would not be misleading to describe the present work as the result of combining Barker's, Kennedy's, and Gibbard's ideas in a pot and simmering them over low heat.

Berkeley, California
June 2025

1 Vagueness and Communication

In a September 1896 letter to Giuseppe Peano, Gottlob Frege writes:

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.¹

The passage defends the uncompromising view of vagueness for which Frege is well known. If we seek rigor and want to apply logic, Frege thinks, we must construct an artificial language in which every term has an independent meaning. 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

¹ Gottlob Frege, *Philosophical and Mathematical Correspondence*, ed. Brian McGuinness, trans. Hans Kaal (University of Chicago Press, 1980), 114–15. “Acervus” is the Latin word for heap: Frege is referring to the paradox better known by its Greek name, “sorites.”

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

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 the color of Red Delicious apples falls into the extension of “red” while that of Valencia oranges does 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 Bertrand 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 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.

³ Solomon Feferman, “Why the Programs for New Axioms Need to Be Questioned,” *Bulletin of Symbolic Logic* 6, no. 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.

⁴ Bertrand Russell, “Vagueness,” *Australasian Journal of Philosophy and Psychology* 1, no. 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 University Press, 2018), 233–35.

Over the past 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 are 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 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 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 this book, I want to focus on another question, to which much less attention has been given. Frege alludes to it briefly in his letter to Peano, immediately after the passage with which we began this chapter. Right 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. 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.⁵

If vague words lack independent meanings, as Frege supposes, how do we manage to use vague language to communicate? Frege's answer to this question consists in two rather cryptic sentences. 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 after we can answer it.

Here is the plan. In this chapter, I try to get clearer about the problem vagueness poses for understanding communication, rejecting some attempts to dismiss it. I argue that the problem is due not to the "fuzziness" of vague language, but to its context 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 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 chapter 2, we 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

⁵ Frege, *Correspondence*, 115. I depart from Brian McGuinness's translation in translating "Satz" as "sentence" rather than "proposition."

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 chapter 3, I present my own approach. I 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 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 in 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, nonfigurative 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 the hearer's recognition of the sentence

uttered by the speaker and by the speaker's and hearer's common knowledge of its sense.⁶

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 its 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 figure 1.1).

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 appropriate only 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.

⁶ Compare David Lewis, "Index, Context, and Content," in *Philosophy and Grammar*, ed. Stig Kanger and Sven Öhman (Reidel, 1980), 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 me to be willing and able to tell the truth, will then be in a position to infer whether A or B or ...".




Figure 1.1: The test tubes

- (c) The tallest test tube is test tube A.

Thus, 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 the assistant is making exactly where the line between the tall and the nontall test tubes falls, or whether there is a line at all. Moreover, the inference is a reliable one, one that can give the hearer *knowledge* of what the speaker meant. It does not rely on past experience of a special use (like an idiom) or a guess about what is meant in context.⁷

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 successful 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 of the test tubes she is demonstrating.

⁷ Pace Timothy Williamson, *Vagueness* (Routledge, 1994), 44.

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 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).⁸

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 communicate. 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 chapter 2, 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

⁸ Williamson, *Vagueness*, 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.”

(5) x is tall iff the height of x is greater than n mm.⁹

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 compositionally work out 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 perfectly 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 with respect to vagueness? For example, consider three modes of presentation of a certain Manhattan street corner:

- A. the northeast corner of Amsterdam Avenue and W 106th Street.
- B. the corner with Mama’s Pizzeria.
- C. the corner depicted in figure 1.2.

All three describe the same corner, but they put their possessors in a different epistemic position with respect to that corner. Mode A enables you to find the corner by deploying your knowledge of Manhattan’s street layout. Mode B enables you to find it by asking locals. Mode C enables you to recognize it if you happen to pass by on the other side of the street. 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

⁹ Here and elsewhere in this book, “iff” abbreviates “if and only if.”



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

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 meters 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.¹⁰ Critics of epistemicism have found this hard to accept, but I am not going to

¹⁰ Williamson, *Vagueness*, 205–12.

contest the claim here.¹¹ 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 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 nontransparent 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.¹² 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 and speaker have matching dispositions to apply "tall" 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 that 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

¹¹ See, for example, Vann McGee and Brian McLaughlin, "Review of Timothy Williamson, *Vagueness*," *Linguistics and Philosophy* 21, no. 2 (1998): 221–35; Vann McGee and Brian McLaughlin, "Logical Commitment and Semantic Indeterminacy: A Reply to Williamson," *Linguistics and Philosophy* 27, no. 1 (2004): 123–36.

¹² In the sense of Tyler Burge, "Individualism and the Mental," *Midwest Studies in Philosophy* 4, no. 1 (1979): 73–122.

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: The speaker and hearer regard themselves as talking about the same height property, and they are not disposed to react to small discrepancies in classification by supposing that they simply attach different senses to the word “tall.”¹³ 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: vague modes of presentation can determine a Fregean *Bedeutung*, and 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: she wants him to clean a specific test tube, it is proximal rather than distal (since “this” was used rather than “that”), and 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 the test tubes. Or the assistant may *guess* which test tube was intended, and

¹³ Burge, “Individualism and the Mental,” section IIIc.

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 communication at issue here is that of success in an illocutionary act, not success in a perlocutionary act.¹⁴

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 (1) the speaker must intend to communicate something about a particular object, and (2) 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 in section 1.2 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.

¹⁴ J. L. Austin, *How to Do Things with Words* (Harvard University Press, 1962), 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 of *which* proximal object 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.¹⁵

1.4 The Context 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 chapters.) 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

¹⁵ This is clear from thinking about how the speaker would react to being corrected. As Burge, “Individualism and the Mental,” notes in discussing his famous example, someone who claims to have arthritis in her 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 will *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. The Carnap/Agnew example discussed in David Kaplan, “Dthat,” in *Syntax and Semantics, Vol. 9: Pragmatics*, ed. Peter Cole (Academic Press, 1978), 239, does not show otherwise. In the scenario Kaplan describes, the speaker intends to use “that” to refer *both* to the picture he’s gesturing at (which is, in fact, a picture of Agnew) 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* 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.¹⁶ 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 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

¹⁶ Herman Cappelen and Ernie Lepore, *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism* (Blackwell, 2005), recognize that their semantic minimalism would not be plausible without “speech act pluralism.”

mean by a sentence like (2) or (8), and in normal cases, it is possible for the hearer to work out which of these is the intended message. For example, one might suppose that in the case of (8) the options 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.

The task of discerning which of these four options the speaker intends is a tractable one: in context, it may be clear that the speaker means (a) rather than (b), (c), or (d). 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 the vicinity of option (a), there are many other things the speaker might mean by (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.

In most cases, hearers will not be in a position to tell 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 10-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 a cushion, a square of fabric, the stuffing behind the fabric, or the wall socket behind the couch, one will have to do something extra to indicate this. None of the candidates (a–j) for reference classes enjoys this kind of “default” status.

Moreover, if you ask someone who utters (9) whether they meant the couch or the cushion, you can expect a clear answer, but if you ask someone who utters (8) whether they meant tall for a third-grader in this school, or in this class, or in this city, you may very well be told, “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 preexisting intention, not a new decision made on the spot in response to your question? Here the analogy with “this” breaks down completely. It seems defective to use “this” without having a determinate object in mind,¹⁷ but using “tall” without having chosen between, say, (f) and (g) does not seem infelicitous in the same way.

¹⁷ In the appendix to chapter 3 we discuss a possible counterexample proposed by Jeffrey King. But King would concede that in most cases, uses of “this” in the absence of determinate referential intentions are infelicitous.

What is more, relativity to a reference class does not exhaust the context sensitivity of gradable adjectives.¹⁸ 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.¹⁹

She proposes that, in addition to a reference class, “old” is 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.²⁰

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 figure 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 context 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:

¹⁸ In addition to the works cited in this discussion, the point is made in Christopher Kennedy, “Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives,” *Linguistics and Philosophy* 30, no. 1 (2007): section 2; Mark Richard, *When Truth Gives Out* (Oxford University Press, 2008), chapter 4; and John MacFarlane, “Vagueness as Indecision,” *Aristotelian Society Supplementary Volume* 90, no. 1 (2016): 255–83.

¹⁹ Delia Graff Fara, “Shifting Sands: An Interest-Relative Theory of Vagueness,” *Philosophical Topics* 28, no. 1 (2000): 66.

²⁰ Fara, “Shifting Sands,” 66.

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?²¹

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.)²² And of course, contextual variation in the ranking is much greater for “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. On these assumptions, successful communication would require the speaker and hearer to coordinate on one of the many truth conditions that the

²¹ Keith DeRose, “Gradable Adjectives: A Defence of Pluralism,” *Australasian Journal of Philosophy* 86, no. 1 (2008): 149.

²² 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 that was constructed in secret and raised 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 (“40 Wall Street,” *Wikipedia*, 2020, https://en.wikipedia.org/w/index.php?title=40_Wall_Street&oldid=941342831, accessed February 28, 2020).

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 called their "tolerance."²³ 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—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.²⁴ Suppose that the particular proposition Alison

²³ Crispin Wright, "Language Mastery and the Sorites Paradox," in *Truth and Meaning*, ed. Gareth Evans and John McDowell (Oxford University Press, 1976), 229.

²⁴ Bacon, *Vagueness and Thought*, 118.

means to express by saying “Harry is bald” has an evidential role that assigns a probability of 0.733, conditional on Harry having 233 hairs (and other relevant precise facts). How is a hearer supposed to recognize that she intends this one, and not one that assigns the conditional probability of 0.698? An appeal to anti-individualism seems the only hope, but as argued earlier, 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 Jeffrey King has called “felicitous underspecification.”²⁵ 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 is for (6). King gives a number of examples.²⁶ For example, looking at a group of surfers on the beach, one might say:

(10) 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 be just 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:

(11) 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

²⁵ Jeffrey C. King, “Strong Contextual Felicity and Felicitous Underspecification,” *Philosophy and Phenomenological Research* 97, no. 3 (2018): 631–57.

²⁶ Examples of this sort are not new, of course. See, for example, François Recanati, *Literal Meaning* (Cambridge University Press, 2003), on “here” and “now” or Cappelen and Lepore, *Insensitive Semantics*, on “ready.”

where Sophie is using her own skis, which she made, a speaker might use “Sophie’s skis” without distinguishing among 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 needed to determine a truth condition. If that is right, the Simple Model of Communication must be rejected.²⁷

1.6 Why Not Diagonalize?

Described abstractly, these examples 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” may wonder whether the approach it takes to such cases, commonly called “diagonalization,” can help us here too.²⁸

Stalnaker assumes 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 to be updated after

²⁷ 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,” 651). If one took “communicative aims” to be perlocutionary aims, one might try to explain these cases as perlocutionary successes that are illocutionary failures. If I aim to insult my host, and to achieve this end I say, “That’s ugly!” without indicating which of her many paintings I am talking about, then my illocutionary act has failed (she does not recognize which proposition I am asserting), but my perlocutionary aim may be achieved (she is insulted nonetheless). Understood this way, King’s cases would not challenge the Simple Model of Communication. As I read King, though, his claim is that coordination on a unique value is not required by the *illocutionary* aims of the speakers, and perfect uptake does not require such coordination. This position is a decisive rejection of the Simple Model.

²⁸ See Robert Stalnaker, “Assertion,” in *Syntax and Semantics, Vol. 9: Pragmatics*, ed. P. Cole (Academic Press, 1978); reprinted in Robert Stalnaker, *Context and Content: Essays on Intentionality in Speech and Thought* (Oxford University Press, 1999), 78–95.

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 formulated as follows:

Uniformity

The same proposition is expressed relative to each possible world in the context set.²⁹

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

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

Prior to my assertion, the common ground does not settle who the woman is: given what is being presupposed, there are many candidates. If the woman 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; if she is Tricia Nixon Cox, then the reference of “that” in context is Cox. So, if (12) were being used to assert a singular proposition—the proposition that this woman is either Gabor or Anscombe—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, non-singular proposition: that the woman who is talking in the next room (whoever she might be) is either Zsa Zsa Gabor or Elizabeth Anscombe.³⁰ This proposition, called the *diagonal*, is true at a world w just in case the proposition that (12) expresses if w is the actual world is true at w (see figure 1.3).

It might seem that this strategy can be applied straightforwardly to all cases of contextual underspecification. In all these cases, we lack common knowledge of the extension of a context-sensitive element, and hence of which proposition is literally expressed by the sentence in question. So

²⁹ Stalnaker, *Context and Content*, 88.

³⁰ This proposition will have a different modal profile than the singular propositions. For example, suppose that the woman talking in the next room is actually Gabor. The singular proposition that Gabor is either Gabor or Anscombe is a necessary truth: it would be true even if the woman talking in the next room were Tricia Nixon Cox. The proposition that the woman talking in the next room is either Gabor or Anscombe, by contrast, is contingent: it would not be true if the woman talking in the next room were Tricia Nixon Cox.

	<i>g</i>	<i>a</i>	<i>c</i>		<i>g</i>	<i>a</i>	<i>c</i>
<i>g</i>	T	T	T	⇒			
<i>a</i>	T	T	T		T	T	F
<i>c</i>	F	F	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 Nixon Cox is the woman speaking. In the “propositional concept” on the left, the rows represent different possibilities for the actual world, while the columns represent different worlds of evaluation. On the right is the diagonal proposition derived from this propositional concept.

why not take the speaker to be asserting the diagonal proposition? In our test tube example, for instance, the diagonal of (2) would be

- (13) 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. 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

- (14) 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 even having a particular referent in mind. This should be just as felicitous as asserting (12) when one does not know whether "that woman" refers to Gabor, Anscombe, or Cox. But it is not. Communication fails in the case of (6) in a way it does not in the case of (12). 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 (12), the speaker and hearer understand 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 figure 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 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 (10), 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. In Irene Heim and Angelika Kratzer’s widely used semantics textbook, deictic pronouns like “she” are taken to be free variables and given their referents by the variable assignment of the context, g_c .³¹ 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),”³² and they add that a context is “appropriate for” a sentence only if it determines values for all the free variables in the sentence. We will be tempted to think we can always diagonalize only if we forget that, when a speaker lacks the requisite intentions, a context will fail to determine a variable assignment.

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 between 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

³¹ Irene Heim and Angelika Kratzer, *Semantics in Generative Grammar* (Blackwell, 1998), section 9.1.2.

³² Heim and Kratzer, *Semantics in Generative Grammar*, 243, emphasis added.

interest us, the presupposition that must be met 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 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 not 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 Simple 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 rather 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 underspecification.” 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 chapters 2 and 3 we will explore how the more general problem might be solved by rethinking the Simple Model of Communication.

In chapter 1, 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 figure 2.1, repeated from chapter 1). 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 Jeffrey King has called




Figure 2.1: The test tubes

“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

- (a) Those guys = Abe, Bob, Cindy, Maria, and Zeke, or
- (b) Those guys = Abe, Bob, Cindy, Maria, and 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 just two, in this case). Our task in this chapter is to explore this idea, which has been endorsed recently by a number of philosophers.¹ I will call the general idea, which different

¹ David Braun and Theodore Sider, “Vague, so Untrue,” *Noûs* 41, no. 2 (2007): 133–56; Ray Buchanan, “A Puzzle about Meaning and Communication,” *Noûs* 44, no. 2 (2010): 340–71; Kai von Fintel and Anthony S. Gillies, “‘Might’ Made Right,” in *Epistemic Modality*, ed. Andy Egan and Brian Weatherson (Oxford University Press, 2011); Jeffrey C. King, “The Metasemantics of Contextual Sensitivity,” in *Metasemantics*, ed. Alexis Burgess and Brett Sherman (Oxford University Press, 2014), 106; Justin Khoo and Joshua Knobe, “Moral Disagreement and Moral Semantics,” *Noûs* 52, no. 1 (2018): 109–43; Jussi Suikkanen, “Contextualism, Moral Disagreement, and Proposition Clouds,” ed. Russ Shafer-Landau, *Oxford Studies in Metaethics* 14 (2019): 47–69; Michael Caie, “Semantic Indecision,” *Philosophical Perspectives* 32, no. 1 (2018): 108–43.

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, there is a view about communication:

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 be surmounted only 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 chapter 3.

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. Their defenders present cases in which the speaker’s intentions do not single out one of a number of possible assertable 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 focus on vague expressions.² 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 appeal to clouds of propositions is motivated not by lexical vagueness, but by cases of felicitous underspecification.⁴ 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). Buchanan suggests, plausibly, that Chet’s intentions may not distinguish between the following options:

- (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.
- (8) Every beer in the apartment is in the bucket.

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).

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

² Braun and Sider, “Vague, so Untrue.”

³ Braun and Sider, “Vague, so Untrue,” 135.

⁴ Buchanan, “A Puzzle about Meaning and Communication.”

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:

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 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.⁷

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.⁸

⁵ Buchanan, "A Puzzle about Meaning and Communication," 350, renumbering the example.

⁶ Buchanan, "A Puzzle about Meaning and Communication," 358.

⁷ See von Fintel and Gillies, "Might' Made Right," 118, renumbering the example.

⁸ See von Fintel and Gillies, "Might' Made Right," 118.

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.⁹ 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 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 (Billy’s) own information. 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:

⁹ See Huw Price, “Does ‘Probably’ Modify Sense?” *Australasian Journal of Philosophy* 61, no. 4 (1983): 396–408; Andy Egan, John Hawthorne, and Brian Weatherson, “Epistemic Modals in Context,” in *Contextualism in Philosophy*, ed. Gerhard Preyer and Georg Peter (Oxford University Press, 2005); Seth Yalcin, “Epistemic Modals,” *Mind* 116 (2007): 983–1026; John MacFarlane, “Epistemic Modals Are Assessment Sensitive,” in *Epistemic Modality*, ed. Andy Egan and Brian Weatherson (Oxford University Press, 2011); John MacFarlane, *Assessment Sensitivity: Relative Truth and Its Applications* (Oxford University Press, 2014), chapter 10. The example in (11) comes from von Fintel and Gillies, “‘Might’ Made Right,” but it is similar to examples used by these relativist and expressivist critics of contextualism.

(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, but von Fintel and Gillies think that we can 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 $\text{Poss}_i K$. Instead, in uttering (11) Alex expresses a *cloud* of propositions, including $\text{Poss}_A K$, $\text{Poss}_B K$, and $\text{Poss}_{AB} K$. Alex is warranted in making her claim because she has warrant for one of the propositions in the cloud ($\text{Poss}_A K$). But Billy's response can target one of the other propositions ($\text{Poss}_{AB} K$).

Although von Fintel and Gillies are mainly concerned with epistemic modals, the way they motivate the cloud-of-propositions picture suggests that they take it to be an appropriate response to contextual underspecification 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 it could not be used in other cases of underdetermination, such as implicit quantifier domains.¹⁰

¹⁰ See von Fintel and Gillies, "Might' Made Right," 123.

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, or asks about them, for example. 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. As we shall see, the three views we are considering give very different answers to this question.

Right away, we face a terminological issue. The term “proposition” is sometimes stipulated to mean the content of an assertion.¹¹ 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 some-things, 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.

There are two ways around this terminological impasse. We can say that the speech acts that 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.

We can find the first approach in von Fintel and Gillies, who talk of speakers “putting in play” clouds of propositions in situations where they would not be entitled to “flat-out assert” all of them.¹² That looks like a denial that the speech act is one of assertion.¹³ 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 so 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.

¹¹ Richard Cartwright, “Propositions,” in *Analytical Philosophy*, ed. R. J. Butler, vol. 1 (Blackwell, 1962).

¹² See von Fintel and Gillies, “‘Might’ Made Right,” 119–20.

¹³ See also von Fintel and Gillies, “‘Might’ Made Right,” 117n18.

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 “determine truth-conditions.”¹⁴ 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 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 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.¹⁵

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 nonborderline 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.”¹⁶ However,

¹⁴ Buchanan, “A Puzzle about Meaning and Communication,” 341.

¹⁵ Braun and Sider, “Vague, so Untrue.”

¹⁶ Braun and Sider, “Vague, so Untrue,” 135.

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.¹⁷

To say that a vague sentence is *approximately true* is to say that all 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 on any legitimate way of disambiguating “bald,” it will express a true proposition.

Thus, when we are ignoring vagueness (as we usually do in everyday life), the norm for asserting or “putting forward” a cloud of propositions is that each proposition in the cloud be true.¹⁸ 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 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.¹⁹

But this simple story is not plausible: The norm it proposes for vague assertion is too stringent. As Chris Barker 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.”²⁰ 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

¹⁷ Braun and Sider, “Vague, so Untrue,” 135.

¹⁸ When we are not ignoring vagueness, on the other hand, *no* assertion of a cloud of propositions can meet the standard for assertion.

¹⁹ Kit Fine, “Vagueness, Truth and Logic,” *Synthese* 30, nos. 3–4 (1975): 265–300; see Caie, “Semantic Indecision,” for an argument that this multi-proposition view is the only plausible way to interpret supervaluationism.

²⁰ Chris Barker, “The Dynamics of Vagueness,” *Linguistics and Philosophy* 25, no. 1 (2002): 1–36; Chris Barker, “Clarity and the Grammar of Skepticism,” *Mind and Language* 24, no. 3 (2009): 253–73; Chris Barker, “Negotiating Taste,” *Inquiry* 56, nos. 2–3 (2013): 240–57.

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.

Although Braun and Sider do not consider this objection, it is one they might meet by appealing to accommodation. *Accommodation*, as David 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 to interpret an utterance as “true, or otherwise acceptable.”²¹ For example, suppose the conversation is currently governed by a high “standard of precision,” so 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, 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.²²

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.

²¹ David Lewis, “Scorekeeping in a Language Game,” *Journal of Philosophical Logic* 8, no. 3 (1979): 347.

²² 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 ” (“Scorekeeping in a Language Game,” 347).

By appealing to this mechanism, Braun and Sider could explain how an assertion of (16) can have the effect of contracting the range 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 Richard counts as tall on every legitimate disambiguation.

However, this reply works only in the case where Richard’s height is mutually known. In that case, we know how to adjust the range of legitimate disambiguations of “tall” in such a way that Richard counts as “tall” on all of them. But suppose Richard’s height is not mutually known. Richard is not in the room with us. We have all seen him before, but we are unsure about his exact height: so 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 being approximately true. 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 than 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 185–193 cm, so 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, whenever 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. David Lewis 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.²³ 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).²⁴

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 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).

²³ Lewis, “Scorekeeping in a Language Game,” 349.

²⁴ 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 coordinate 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.”

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? It can’t be an update to the nonfactual scoreboard, because the second disjunct of (19) is compatible with any threshold for “tall.” And it can’t be 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). Looking at a group of surfers on the beach, we say:

- (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) Abe, Bob, Cindy, Maria, and Zeke, or
- (b) Abe, Bob, Cindy, Maria, and 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’s 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." 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."²⁵ 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."²⁶ Thus,

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.²⁷

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

²⁵ Buchanan, "A Puzzle about Meaning and Communication," 353.

²⁶ Buchanan, "A Puzzle about Meaning and Communication," 366n22.

²⁷ Buchanan, "A Puzzle about Meaning and Communication," 359.

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 nonequivalent, yet equally correct ways.”²⁸

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.²⁹ 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 to be a peripheral part of Paul 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 of Grice’s account: the idea that the speaker’s meaning intention is *inter alia* an intention that the audience recognize this intention.³⁰ 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 observes, it is the transparency of Gricean meaning intentions that allows us to think of speech acts as updating a common ground.³¹ 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.

²⁸ Buchanan, “A Puzzle about Meaning and Communication,” 359.

²⁹ Unless it is already common ground that they will serve only the beers they bought at the bodega.

³⁰ Buchanan seems to recognize that he needs to modify this part of Grice’s view (“A Puzzle about Meaning and Communication,” 368n34).

³¹ Robert Stalnaker, “Common Ground,” *Linguistics and Philosophy* 25, nos. 5–6 (2002): 704.

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?³² 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 onto. 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 common 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 responses 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 gives good results only in the specific case of epistemic possibility claims.

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

- (22) It might be that *P*

³² Perhaps Buchanan resists this approach on the grounds that the same worries about determinacy that arise for propositions can be raised for proposition types. That is, there are many 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.

“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.³³

Two exegetical notes, before we continue. First, although von Fintel and Gillies formulate 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’s wording suggests only a necessary condition:

Suppose an utterance of *might*(B)(ϕ) by S puts in play the propositions $P_1, P_2, \dots$. 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.³⁴

³³ See von Fintel and Gillies, “‘Might’ Made Right,” 120.

³⁴ See von Fintel and Gillies, “‘Might’ Made Right,” 120.

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.”³⁵ 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 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 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.³⁶

³⁵ See von Fintel and Gillies, “‘Might’ Made Right,” 120, emphasis added.

³⁶ See von Fintel and Gillies, “‘Might’ Made Right,” 121.

This constraint on the hearer's choices is articulated in the following principle:

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).³⁷

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 $\text{Poss}_A K$: 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, she is targeting $\text{Poss}_{AB} K$, while if she *rejects* it, she is targeting $\text{Poss}_B K$.

If we are going to take into account all 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

³⁷ See von Fintel and Gillies, "Might' Made Right," 121. As with Assert, I have restated this as a general principle about cloudy speech acts, rather than bare epistemic modal claims.

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 contributions. 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's 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 $\neg \text{Poss}_{AB} K$.³⁸ On the other hand, if Billy accepts Alex's claim, then it becomes common ground that $\text{Poss}_{AB} K$ 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."³⁹

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's 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.

³⁸ See von Fintel and Gillies, "‘Might’ Made Right," 123–24.

³⁹ See von Fintel and Gillies, "‘Might’ Made Right," 124.

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 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 keys 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 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.⁴⁰ 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 works there only because of a special property of these modals that is not even shared by epistemic modals in general.⁴¹

⁴⁰ Adding information might, depending on the case, make something either more probable or less probable. For example, learning that a horse has won all its previous races will increase our subjective probability that it will win this race, while 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 $\text{Prob}_{AB} K$.

⁴¹ 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 .

In this particular case, one might try to salvage things by (1) arguing that Billy is not in a position to have a reasonable opinion about $\text{Prob}_{AB} K$, and (2) 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 that 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 long as Alex knows all of this, and knows that Billy has no 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, Billy's information may be strictly stronger than hers. That is, for all Alex knows, Billy may know everything that Alex knows that is relevant to the location of the keys, plus more, and Billy may know 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.
- (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

This is what gives us the logical relations between $\text{Poss}_A K$ and $\text{Poss}_{AB} K$ that we need for von Fintel and Gillies's account to work, and what is missing with Prob.

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. In this scenario, Billy 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 be relevant only 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 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. 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 gradable adjectives. Suppose Anna sees Tim, who stands 1,900 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 1,850 and 2,000. 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 1,850 mm tall, since that is already common ground. What about 1,851 mm? 1,852 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 underspecification. It is presented by von Fintel and Gillies as a conservative alternative to heretical relativist and expressivist views. And it *is* conservative in one respect: it preserves 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 proposals for what this new picture could look like and found them all 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 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 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’s 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 rather 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 non-factual parameter. He describes ruling out delineations as reducing our "ignorance" and "uncertainty" about features of our discourse.⁴² 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.⁴³

On this interpretation, Barker's view reduces to the combination of epistemicism and diagonalization that we considered at the end of chapter 1 (see figure 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" determined by the "discourse" of that 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 chapter 1, diagonalization is possible only when 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 should find it perfectly appropriate, when we know Richard's exact height but consider him

⁴² Barker, "The Dynamics of Vagueness," 3–4, 9.

⁴³ Barker, "The Dynamics of Vagueness," 5–6; for similar remarks, see Barker, "Clarity and the Grammar of Skepticism," 257; Barker, "Negotiating Taste," 246.

	<i>a</i>	<i>b</i>	<i>c</i>		<i>a</i>	<i>b</i>	<i>c</i>
<i>a</i>	$h_a \geq t_a$	$h_b \geq t_a$	$h_c \geq t_a$	$\Rightarrow$			
<i>b</i>	$h_a \geq t_b$	$h_b \geq t_b$	$h_c \geq t_b$		$h_a \geq t_a$	$h_b \geq t_b$	$h_c \geq t_c$
<i>c</i>	$h_a \geq t_c$	$h_b \geq t_c$	$h_c \geq t_c$				

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 . As in figure 1.3, the rows represent different possibilities for the actual world, while the columns represent different worlds of evaluation. 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 .

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) seem inappropriate (these sentences would normally be taken to imply 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. 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.⁴⁴ 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 chapter 3.

⁴⁴ Allan Gibbard, *Thinking How to Live* (Harvard University Press, 2003).

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 Robert Stalnaker, we can think of the assertion as a proposal to modify the common ground in some way.¹ Understanding it (*uptake*) consists in recognizing how the speaker proposes to update the common ground. But this is schematic: To fill out the story, we need to say more about what the common ground consists in, and how an assertion of (1) proposes to update it. That has been our concern in chapters 1–2.

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 chapter 1, (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.

¹ Robert Stalnaker, "Assertion," in *Syntax and Semantics, Vol. 9: Pragmatics*, ed. P. Cole (Academic Press, 1978).

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 chapter 2, we considered some answers to these questions from the literature and found them unsatisfactory. On Ray Buchanan's view, uptake is achieved no matter which proposition in the cloud the hearer fastens onto.² 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.³ 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 David Braun and Theodore Sider, take the speaker to be proposing that *all* the propositions in the cloud be added to the common ground.⁴ This approach yields a version of supervaluationism. But, as we saw in chapter 2, 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.⁵ Someone who asserts

² Ray Buchanan, "A Puzzle about Meaning and Communication," *Noûs* 44, no. 2 (2010): 340–71.

³ See Kai von Fintel and Anthony S. Gillies, "Might' Made Right," in *Epistemic Modality*, ed. Andy Egan and Brian Weatherson (Oxford University Press, 2011).

⁴ David Braun and Theodore Sider, "Vague, so Untrue," *Noûs* 41, no. 2 (2007): 133–56.

⁵ Following Chris Barker, "The Dynamics of Vagueness," *Linguistics and Philosophy* 25, no. 1 (2002): 1–36, and David Lewis, "General Semantics," *Synthese* 22, nos. 1–2 (1970): 18–67, I define a *delineation* as a function from gradable adjective meanings to thresholds. The *s* function in Christopher Kennedy, "Vagueness and Grammar: The

(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 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's pioneering work on the dynamics of vague discourse.⁶ But how is this formal representation of the common ground to be understood? Barker's approach is to think of a delineation as a special kind of fact—a fact about the "discourse."⁷ If that is right, then accepting a content represented as a set of world/delineation pairs is just having a factual belief. This leads Barker to conceive of assertions as reducing our "ignorance" or "uncertainty" about "the" threshold governing the discourse.⁸ The epistemicist's ideology has returned. Indeed, so understood, Barker's view is exactly the view considered at the end of chapter 1, which combines the epistemicist view that the context determines a threshold with diagonalization.

Semantics of Relative and Absolute Gradable Adjectives," *Linguistics and Philosophy* 30, no. 1 (2007): 1–45, is a delineation in this sense, as Kennedy observes, citing Lewis. But where Kennedy supposes that a context provides a delineation, Lewis does not. Lewis says explicitly that "the delineation coordinate is noncontextual. It resembles the assignment coordinate, in that we will ordinarily generalize over it rather than hold it fixed" (Lewis, "General Semantics," 229).

⁶ Barker, "The Dynamics of Vagueness."

⁷ Barker, "The Dynamics of Vagueness," 5–6. Formally, the delineation is determined as a function of the world of evaluation.

⁸ Barker, "The Dynamics of Vagueness," 3–4, 9.

We rejected that approach on the ground 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 rather 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/delineation pairs to represent the content of an attitude, or the common ground of a conversation? That is the question we will attempt to answer in this chapter. 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 Timothy Williamson argue that we have inexact knowledge of the locations of thresholds for vague gradable adjectives, and it is impossible to know where these thresholds lie because knowledge requires a margin for error.⁹ 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 the following:

- (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.

⁹ Timothy Williamson, *Vagueness* (Routledge, 1994).

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:

- (6) ??Miguel is probably tall.
- (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.¹⁰ 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.

There is a clear contrast here with ordinary cases of inexact knowledge. We would 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.

Or, if we did say these things, we would be understood to be conveying that some vagueness in “taller than” affects whether Miguel counts as taller than the pole.¹¹ We would not be understood to be conveying uncertainty about the pole’s height.

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.

¹⁰ For this terminology, see Stephen Schiffer, *The Things We Mean* (Oxford University Press, 2003), chapter 5.

¹¹ For example, when determining whether Miguel is taller than the pole, do we measure him with or without shoes?

3.2 Plan Expressivism

The fundamental idea of 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 section 3.1—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.¹² *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 way 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, $D \vee P$. 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, and 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

¹² Allan Gibbard, *Thinking How to Live* (Harvard University Press, 2003).

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*.

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.¹³

¹³ 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," for example, the semantic values of gradable adjectives are degree functions, and phrases like "for an athlete" modify these functions by restricting their domains.

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 chapter 1 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 these things as being determined by a hyperplan. After all, a hyperplan is a maximally specific plan, one that includes plans for *everything*—including, a fortiori, these aspects of the use of language. Although Barker was wrong to suggest that a possible world determines a delineation, a hyperplan does determine a delineation.

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:

(16) $\{(w, h) \mid \text{Richard's height in } w \geq \text{the threshold for "tall" set by } h\}$.¹⁴

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.

¹⁴ I ignore complexities introduced by reference classes (see note 13). A fuller account would talk of the threshold for *tall* for an *F* determined by *h*, where *F* is the reference class determined by *h* for this occurrence of “tall.”

- 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 that thresholds (or threshold ranges) governing vague words can shift as a conversation unfolds.¹⁵ 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 expressivist view stand on these traditional questions? That depends on how we formulate bivalence and 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. In general, a context does not 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.

¹⁵ See, for example, Delia Graff Fara, “Shifting Sands: An Interest-Relative Theory of Vagueness,” *Philosophical Topics* 28, no. 1 (2000): 45–81; Scott Soames, *Understanding Truth* (Oxford University Press, 1999).

But there are other formulations of bivalence that are compatible with expressivist ideology. Following Williamson, we can define monadic truth and falsity predicates, applicable to utterances, as follows:¹⁶

- (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.¹⁷ Using these predicates, we can formulate bivalence as a schema:

Bivalence 2

- If u is an utterance that says that ϕ , then u is either true or false.¹⁸

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

- (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.¹⁹ 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

- (7) 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.

¹⁶ Williamson, *Vagueness*, 188.

¹⁷ Williamson, *Vagueness*, 192.

¹⁸ 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).

¹⁹ More properly, each gradable adjective meaning (allowing for sensitivity to a reference class): see note 13.

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:

- (19) 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 *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 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 nontall, just as we can be undecided about whether to buy Crest or Colgate toothpaste. But could we not also plan *not* to draw a sharp line between the tall and the nontall—just as we might decide to be indifferent between Crest and Colgate? A decided but *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).²⁰

²⁰ Allowing this kind of indifference in plans amounts to rejecting the natural idea that a fully decided plan should say *what to do*: A plan that is indifferent between several

Gibbard's view allows for fully decided plans that are indifferent between options in this way. A Gibbardian hyperplan maps each circumstance, not to an action, but to a nonempty *set* of permissible actions, between which it is indifferent.²¹ This allows Gibbard to distinguish the *undecided* state of mind:

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

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

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

The content of (20) contains hyperplans that map the present circumstance to { *buy Crest* } and hyperplans that map it to { *buy Colgate* }, while the content of (21) 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" and distinguish between (20) and (21).²² But it is also potentially problematic for his aim of giving a naturalistic account of normative states of mind. 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

options does not do this. To be sure, a plan might say to buy Colgate if a coin lands heads and Crest if it lands tails. But is that the same as the plan's being indifferent between the options of buying Colgate and buying Crest? Or is it rather a plan to perform a third action that is incompatible with both? Consider the spectrum of actions of the form *role a die and buy Colgate if it lands on n or less, otherwise buy Crest*. When n is 1 or 5, such actions certainly do not seem to reflect indifference between buying Colgate and buying Crest. It is not clear that the situation is any different when n is 3.

²¹ Gibbard, *Thinking How to Live*, 56.

²² Gibbard, *Thinking How to Live*, 54–56; James Dreier, "Relativism (and Expressivism) and the Problem of Disagreement," *Philosophical Perspectives* 23, no. 1 (2009): 79–110; Mark Schroeder, *Being for: Evaluating the Semantic Program of Expressivism* (Oxford University Press, 2008).

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.²³

Here, 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.²⁴ So nothing about the use to which we want to put plan expressivism compels us to countenance hyperplans that are indifferent among several candidate thresholds. For our purposes, we could think of hyperplans as the contents of fully determinate conditional intentions, which select a unique action for each possible circumstance.

On the other hand, nothing prevents us from allowing hyperplans to be indifferent among 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, 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 logic) that rejects the principle of 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

$$\begin{aligned} (22) \quad f(w, h) &= 1 && \text{if } \forall t \in T_h \, r_w \geq t \\ &= 0 && \text{if } \forall t \in T_h \, r_w < t \\ &= \frac{1}{2} && \text{otherwise} \end{aligned}$$

where T_h is the set of thresholds for “tall” permitted by h and r_w is Richard’s height in w . An expressivist who develops the view in this way

²³ I am indebted here to an insightful unpublished paper, “Preference and the New Negation Problem,” by Sophia Dandelet.

²⁴ The negation problem arises because we need to distinguish three states of mind, which we can describe by putting negations in different places: (1) not believing one ought to ϕ , (2) believing it is not the case that one ought to ϕ , and (3) believing one ought not to ϕ . Without indifferent hyperplans, the expressivist only has resources to describe two of these: (1) not planning to ϕ and (3) planning not to ϕ . But with gradable adjectives, we do not have so many places to put negations. We only need to distinguish between (1) not believing that a six-foot-tall man is tall and (2) believing that a six-foot-tall man is not tall. (1) corresponds to not planning to count a six-foot-tall man as tall, and (2) corresponds to planning to count a six-foot-tall man as tall.

would reject the principle of excluded middle, and hence Bivalence 2. However, this motivation for rejecting classical logic is very different from the usual motivations: accounting for borderline cases and solving the sorites paradox. Nonclassicality would be called for because we can be decidedly indifferent about where the threshold lies, not just undecided about it. There may be reasons to countenance this kind of indifference, but it is not necessary for accounting for borderline cases or the sorites paradox. As I will argue shortly, indecision is enough for that.

3.4 The Sorites Paradox

Up to this point, 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:

- (23) If people whose net worth is n dollars are rich, so are people whose net worth is $n - 1$ dollars.
- (24) People whose net worth is a billion dollars are rich.
- (25) 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 accept classical logic, this means accepting the negation of (23):

- (26) 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 paradox. But accepting (26) is less problematic for the expressivist than for the supervaluationist. A supervaluationist takes (26) 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 (26) is accepted in a common ground is just to say that all the hyperplans that we are undecided among 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.... Similarly, a group may plan to go *somewhere* for dinner without there being any particular restaurant where they plan to go. Such is indecision.

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 (23) so hard to reject, if its negation is accepted in the common ground?

The attractiveness of (23) 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.²⁵ 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 (23).

As Diana Raffman, Scott Soames, Delia Graff Fara, and others have pointed out, this is bad reasoning.²⁶ There may be an explanation for our inability to conceive of a one-dollar boundary between the rich and the nonrich 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 (23), without giving us any reason to suppose that (23) is true.

I agree that the appeal of (23) rests on a false inference from the difficulty of thinking of a counterexample. However, 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.²⁷ 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

²⁵ Crispin Wright, "Language Mastery and the Sorites Paradox," in *Truth and Meaning*, ed. Gareth Evans and John McDowell (Oxford University Press, 1976).

²⁶ Diana Raffman, "Vagueness Without Paradox," *Philosophical Review* 103, no. 1 (1994): 41–74; Fara, "Shifting Sands"; Soames, *Understanding Truth*.

²⁷ 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* 65, no. 2 (2002): 422–28, rightly asks what we have gained by moving from one hidden boundary to another.

“rich” prevents this, and we can imagine cases where it might be useful to adopt such a plan.²⁸ I gave my first talk on this subject in Cardiff, in the shadow of The Garth. This hill was the inspiration for the 1995 film *The Englishman Who Went Up a Hill but Came Down a Mountain*, which tells the story of Welsh villagers who, dismayed at the reclassification of their local mountain as a hill, carried buckets of earth to the top to push it over the 1,000-foot threshold for counting as a mountain in official maps of Britain at the time.²⁹ High points 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 nonrich, how can we explain our intuitive repugnance to instances of (26)? On the view we have been exploring, accepting an instance of (26) 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

²⁸ The claim about meaning is controversial. Jerry Fodor and Ernie Lepore, “What Cannot Be Evaluated Cannot Be Evaluated and It Cannot Be Supervalue Either,” *Journal of Philosophy* 93, no. 10 (1996): 516–35, argue 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” (Kennedy, “Vagueness and Grammar,” section 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.

²⁹ *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). Although the film is about a fictional hill called “Ffynnon Garw,” its writer and director was born in Taff’s Well, underneath The Garth, and reportedly heard the story from his grandfather: Ed Sullivan, “Movie Review: Ffynnon Garw,” *Professional Surveyor* 18, no. 8 (1998), <https://archives.profsurv.com/magazine/article.aspx?i=368>.

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 toward accepting it.³⁰

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 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.³¹ Of course, if I am faced with a sequence of such proposals, I will eventually have to reject one—or else 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.³² And if I am right that vague assertions are partly proposal-like, the two cases can be explained in exactly the same way.

³⁰ See John MacFarlane, “Vagueness as Indecision,” *Aristotelian Society Supplementary Volume* 90, no. 1 (2016): 272–74.

³¹ It is true that if we imagine specific sorts of situations—perhaps balloons come in packs of 5, 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 previously) where it makes sense to agree on a sharp boundary.

³² See Diana Raffman, *Unruly Words: A Study of Vague Language* (Oxford University Press, 2014), chapter 5; Raffman, “Vagueness Without Paradox.”

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 nontall inevitably seem to reintroduce sharp boundaries elsewhere.³³ For example, supervaluationists avoid commitment to an inscrutable threshold by considering a range of admissible thresholds, but the end points 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?

On the expressivist view, as we have seen, the contents asserted do not depend on the range of admissible thresholds. Thus, 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.³⁴ If someone asserts (2), 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 end up with common knowledge of the *result* of the update only if we had common knowledge of the *prior* state of the context. Before (2) 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 (2) that the shortest person who had been called "tall" in the conversation was 195 cm, the tallest person who had been called "not tall" was 180 cm, and 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

³³ Williamson, *Vagueness*.

³⁴ Or true common belief, in the version of Robert Stalnaker, "Common Ground," *Linguistics and Philosophy* 25, nos. 5–6 (2002): 716. The difference between common belief and common knowledge does not matter for the arguments I make here.

$180 < t \leq 195$. But what about a case in which (2) 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 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.³⁵ This is an instance of the principle:

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 $x - \delta$ counts as tall. But common knowledge does not satisfy the principle:

Negative Introspection

If it is not common knowledge that p , it is common knowledge that it is not common knowledge that p .³⁶

³⁵ It is common knowledge that p in a group just in case each member knows that p , knows that each member knows that p , knows that each member knows that each member knows that p , and so on *ad infinitum*.

³⁶ See Stalnaker, “Common Ground,” 707.

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.

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 Egypt is in the northern hemisphere, or that Ohio is east of Illinois, or that there are two square roots of 4, or that Michelangelo painted the Sistine Chapel ceiling, or that Richard Nixon resigned the US presidency? One can start asking hard questions about where, exactly, common ground leaves off. However, an inability to answer these questions does not show that there is not common ground. Because it is not in general 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 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 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? It is common for philosophers and cognitive scientists to suppose that fish and frogs can have mental representations with contents like *that is food*, *there is danger*, or *it is cold here*. Are these attributions unwarranted unless we are willing to ascribe to fish and frogs the ability to deliberate about where to draw the line between food and nonfood, between danger and nondanger, between cold and noncold?

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. That is not to say that they are precise. Rather, the distinction between vagueness and precision is one that we can sensibly draw only for creatures that can evince the sort of evidence brought out by sentences like (3–11) in section 3.1. We

can use vague sentences to ascribe thoughts to primitive creatures, but we need to keep in mind the limitations of doing this.³⁷

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 and practical.

This account has the resources to answer an important question raised by Stephen Schiffer: What, exactly, is it to *take something to be borderline*?³⁸ 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 nontall 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 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 less 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 that our ambivalence in borderline cases is not doxastic uncertainty.³⁹

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

³⁷ A bolder way of making this point is that epistemicism might be just fine as an account of the frog's *food* thoughts. But that does not show that it is fine as an account of *our food* thoughts.

³⁸ Stephen Schiffer, "Two Issues of Vagueness," *The Monist* 81, no. 2 (1998): 193–214; Schiffer, *The Things We Mean*, chapter 5.

³⁹ This argument cuts not only against classical epistemicism, but against views like that of Andrew Bacon, *Vagueness and Thought* (Oxford University Press, 2018), 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*," *Philosophical Review* 129, no. 1 (2020): 153–58.

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 50–50 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 as 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 previously, toward the conjunctive proposition that Miguel is tall, bald, smart, *and* funny.

These contrasts are all plausible, but they serve to pick out a subject matter rather 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 have any coherent account of mixed states of mind, which involve both sources of partiality (vagueness and uncertainty). 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.⁴⁰

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 both vagueness and uncertainty. On this view, one's attitude toward propositions that 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 because 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 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.⁴¹ There is no way of collapsing this range to a single number.

⁴⁰ John MacFarlane, "The Things We (Sorta Kinda) Believe," *Philosophy and Phenomenological Research* 73, no. 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 University Press, 2010).

⁴¹ If probabilities are defined for sets of worlds, then we can derive a function from hyperplans to probabilities: $f(h) = Pr(\{w | \langle w, h \rangle \in P\})$, where $P \subset W \times H$ is a Gibbardian content. From this function, we can derive a set of probabilities (the function's range), or an upper and lower bound. But both moves lose information.

One might think we can assign probabilities to our various possible threshold decisions and then 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 of how likely it is that one will plan to put a threshold in a particular place with the question of how likely it is that the threshold *is* in that place. There is no making sense of the latter, so long as the issue is left undecided.⁴²

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 man 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

⁴² We can illustrate this point by considering a similar move that 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 here about which we might have a credence of 0.8. She has only conditional preferences about where to live, not unconditional ones.

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 partial ordering to a quantitative measure.

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* in 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."⁴³ In these cases, we can safely act as if the probability is somewhere in that range, so long as our actions do not depend on *where* it is in the range.

Indeed, it is sometimes possible to assign sharp probabilities to vague sentences because indecision about the threshold matters only for states of affairs we have already ruled out. For example, if our plans constrain

⁴³ See, for example, Isaac Levi, "On Indeterminate Probabilities," *Journal of Philosophy* 71, no. 13 (1974): 391–418; James M. Joyce, "A Defense of Imprecise Credences in Inference and Decision Making," *Philosophical Perspectives* 24, no. 1 (2010): 281–323; and Susanna Rinard, "A Decision Theory for Imprecise Probabilities," *Philosophers' Imprint* 15, no. 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.

the threshold for counting as tall to between 185 and 195 cm, and we have a 0.5 credence that Miguel is exactly 180 cm in height and a 0.5 credence that he is exactly 200 cm in height, 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 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 to handle the similar issue of judgments with mixed normative and descriptive content.

3.7 Why Be Vague?

I want to close this chapter by discussing a question that I do not think is asked enough: 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”? An account of vagueness ought to give us resources to answer this question.

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 “tall” is useful when we do not know the heights of things precisely. Even when we do not know a thing’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 tall. 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 (contextually determined) reference class”?

I want to propose that the expressivist view developed here offers a compelling solution to this problem. In a nutshell, vague words are useful because they facilitate joint planning. Our joint endeavors often require us to draw lines and make distinctions. For example, when hiring people, we need to distinguish the qualified from the unqualified, and when designing filtering software, we need to distinguish the pornographic from the nonpornographic. Getting agreement about where the lines should be drawn 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, as well as 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 the 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 about 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.⁴⁴

Inquiry would be impoverished if we heeded Voltaire’s advice: “Define your terms, I say, or we will never come to agreement.”⁴⁵ 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. It’s often better to leave lines undecided while we achieve consensus on the constraints governing them and the consequences of drawing them.

⁴⁴ For kindred thoughts, see Jamie Tappenden, “Some Remarks on Vagueness and a Dynamic Conception of Language,” *Southern Journal of Philosophy* 33, no. supplement (1994): 198.

⁴⁵ “Définissez les termes, vous dis-je, ou jamais nous ne nous entendrons.” Voltaire, *Dictionnaire philosophique*, vol. 4, *Oeuvres Complètes* (Hachette, 1960), 205.

Appendix: On King's Theory

Jeffrey King's valuable book *Felicitous Underspecification*, which appeared a year after the publication of the Woodbridge Lectures, develops a theory of assertoric update that is distinct enough from the views criticized in chapters 1–2 that it deserves a separate discussion.⁴⁶ To preserve the structure of the original lectures on which chapters 1–3 are based, I discuss King's view here, in a separate appendix.

Like Braun and Sider and von Fintel and Gillies, King thinks that in cases of felicitous underspecification, there are multiple candidates for the proposition expressed by an utterance, but the update proposed by the speaker's assertion is the addition of a single proposition to a standard, Stalnakerian common ground. Given these assumptions, a theory of assertoric update must be an account of how the asserted proposition is determined from the set of candidates and other aspects of the conversational context. In chapter 2, we rejected the theories of assertoric update given by Braun and Sider and von Fintel and Gillies. King's theory is different.

On King's view, the update is generally either the conjunction or the disjunction of the candidate propositions. We determine which it is by reasoning pragmatically about how well they would answer the question under discussion (QUD).⁴⁷ More precisely:

Felicitous underspecified update (FUU)

Given c 's context set cs , update cs with the weakest candidate propositional update U for ϕ in c , if any, such that: (1) it gives a partial answer to the immediate question under discussion while adhering to Gricean maxims and not being ruled out by the common ground; and (2) no stronger candidate propositional update for ϕ in c gives a better answer to the immediate question under discussion than does U while adhering to Gricean maxims and not being ruled out by the common ground.⁴⁸

To simplify this a bit, say that proposition is *in the running* to be the update if the following are all true:

⁴⁶ Jeffrey C. King, *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context* (Oxford University Press, 2021).

⁴⁷ See Craige Roberts, "Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics," *Semantics and Pragmatics* 5, no. 6 (2012): 1–69.

⁴⁸ King, *Felicitous Underspecification*, 39.

- It is either the conjunction or the disjunction of the candidate propositions.
- It gives a partial answer to the immediate QUD.
- It is compatible with the common ground.
- It adheres to Gricean maxims.

King's update rule then says: Of those propositions that are in the running to be the update, the update is the one that gives the best answer to the immediate QUD, or if both give equally good answers, it is the weaker one (that is, the disjunction). If no proposition is in the running, the update is not defined, and the assertion is infelicitous.

We can best understand how this works by considering two examples, one where the predicted update is the conjunction of the candidate propositions, and one where it is the disjunction. Here is the first example. We are looking at a car. I say:

(27) That is a beautiful car.

According to King, it need not be determined here whether a *car type* or a *car token* is at issue.⁴⁹ (Am I talking about the Ferrari F50 model in general, or that particular Ferrari F50 in front of us, the one with the blue paint?) As evidence for this, King notes, I could felicitously follow up with either of these sentences, and you would not feel as if I had changed the subject:⁵⁰

(28) Yes, and it was the first Italian commercial sports car to go 0–60 in under 2.5 seconds.

(29) Oh, but unfortunately, it has a scratch on the hood.

So here, we have a small cloud of candidate propositions:

(30) That [*car-token in front of us*] is a beautiful car.

(31) That [*car-type, Ferrari F50*] is a beautiful car.

In this case, King argues, the update is the conjunction of (30) and (31). Let us see why. The QUD, King assumes, is “Is that a beautiful car?” Of course, this question is itself underspecified (is it about a token

⁴⁹ King, *Felicitous Underspecification*, 25.

⁵⁰ King, *Felicitous Underspecification*, 26.

or a type?). King argues that it must be understood conjunctively, as a question about both the car token and the car type.⁵¹ King then argues that the disjunction of (30) and (31) would not settle either question, whereas the conjunction would settle both. Hence, the conjunction is a better answer to the QUD, and nothing about the common ground rules it out, nor does it violate Gricean maxims.⁵² So the update in this case is conjunctive.

The second example involves underdetermination of the threshold for a gradable adjective. At a family gathering, King imagines, he sees his nephew, Colin, and remarks:

(32) Colin is tall.

Here, King says, "It is easy to imagine that, in uttering as I did, I did not intend for a specific degree of height to count as the threshold for being tall in this context. Nor did my audience take me to be intending a specific degree."⁵³ Suppose that it is not common ground exactly how tall Colin is, and that the candidate thresholds are 5'11", 6'0", and 6'1", so that the candidate propositions are:

(33) The degree of Colin's height exceeds 5'11".

(34) The degree of Colin's height exceeds 6'0".

(35) The degree of Colin's height exceeds 6'1".

Then, King argues, the propositional update is the disjunction of (33), (34), and (35), for the QUD is "Is Colin tall?" Of course, this question contains "tall" and is plausibly subject to underspecification. This time King interprets it disjunctively:

Since the candidate semantic values for "tall" in this context are 5'11", 6'0", and 6'1", there is no principled reason to take the question to be whether Colin's height exceeds any one of these thresholds. Hence, we should take the immediate question under discussion to be in effect "Does Colin's height exceed 5'11", 6'0", or 6'1"?"⁵⁴

⁵¹ King, *Felicitous Underspecification*, 46–47.

⁵² King, *Felicitous Underspecification*, 47.

⁵³ King, *Felicitous Underspecification*, 33.

⁵⁴ King, *Felicitous Underspecification*, 48.

The disjunction of (33), (34), and (35) is a complete answer to this question. And stronger updates, such as the conjunction of (33), (34), and (35), are incompatible with the common ground, since it is common ground that 5'11" and 6'0" are candidate thresholds, and this update, King argues, would be incompatible with these thresholds.⁵⁵ Thus, the propositional update in this case is disjunctive.

The central idea here is that the QUD may allow speaker and hearers to arrive at a determinate update to a Stalnakerian common ground, even in the presence of felicitous underspecification. The update will sometimes be the conjunction of the candidate propositions, and sometimes the disjunction, depending on the QUD and other aspects of the context. But I am skeptical that King's approach has any real explanatory power. For in general, the QUD will be underspecified in much the same way as the speakers' assertions. In (27), where the QUD is "Is that car beautiful?," conversational participants may not have settled whether the question is about a type or a token. In (32), where the QUD is "Is Colin tall?," they may not have settled what counts as being tall. King can derive his predictions about the propositional update only *after* assuming a particular determinate understanding of the QUD. In (27), he predicts a conjunctive update after assuming that the QUD should be understood conjunctively. In (32), he predicts a disjunctive update after assuming that the QUD should be understood disjunctively. In both cases, a different resolution of the indeterminacy in the QUD would have led to a different prediction about the propositional update. So, we

⁵⁵ It is not clear to me why King thinks that throwing out worlds where Colin's height is 6'0" amounts to ruling out 6'0" as a candidate threshold for "tall" (King, *Felicitous Underspecification*, 48–49). After all, King doesn't think the *disjunctive* update, which leaves open the possibility that Colin is 6'0", rules out 6'1" as a candidate threshold for "tall." King says: "The fact that conversational participants update with [the disjunction] after hearing ["Colin is tall"] doesn't mean that ["Colin is tall"] expresses [the disjunction], nor that the semantic value in context of "tall" in ["Colin is tall"] is 5'11", so that we no longer have a case of felicitous underspecification" (King, *Felicitous Underspecification*, 35). But then, equally, we should think that the fact that the participants update with the conjunction doesn't mean that "tall" has 6'1" as its semantic value in the context, or that the other candidate semantic values are no longer in play. Note, however, that even if King is wrong about this and the conjunction is compatible with the common ground and thus in the running to be the update, FUU would still select the disjunctive update, because it is weaker but gives an equally good answer to the QUD.

have an account that predicts a propositional update only given a prior decision about how to construe the QUD. But for that, no explanatory account is given.

It is also hard to see what justifies King's claim that the updates stronger than the disjunctive one are ruled out in (32) because they are incompatible with common-ground assumptions about the range of admissible thresholds. Assertoric updates can reduce the set of possible worlds in the common ground: Why should they not also be able to reduce the set of admissible thresholds? Indeed, King himself admits that there are cases where the primary assertoric update is a change in the admissible thresholds. If it is common ground that Colin is 6'0", for example, then the effect of asserting (32) is to rule out thresholds higher than that.⁵⁶ King calls these updates "metalinguistic" and regards them as a special case that is outside the scope of his analysis.⁵⁷

But as we have seen, it is a mistake to think that there are just two kinds of cases: cases where the assertoric update is purely factual and cases where it just affects the admissible thresholds. There are also "mixed cases," in which the update is neither purely factual nor purely nonfactual, nor can it be factored into a combination of factual and nonfactual updates. In these cases, the update can only be understood as imposing a constraint on *combinations* of worlds and thresholds.

Indeed, it is plausible to understand (32) as a case of this kind. Although it isn't common ground exactly how tall Colin is, the speaker may mean to institute the constraint that, going forward, we consider only thresholds that count Colin as tall. If we later find that he is 6'0", for example, we will then be committed to ruling out 6'1" as an admissible threshold. This conditional commitment is lost if we reduce the assertoric update to the disjunction of (33), (34), and (35)—or, equivalently, to (33). The problem with King's theory is that it has no way of representing the kind of assertoric update we get in mixed cases.

We can also understand (27) as a mixed update. The speaker isn't yet committed, in asserting (27), to both the type and the token being beautiful. The conversation could evolve in different ways—either as being

⁵⁶ King, *Felicitous Underspecification*, 34.

⁵⁷ King, *Felicitous Underspecification*, 8, 34, 41.

about the type or as being about the token—and only once this has become determinate will the speaker acquire a specific factual commitment. So, for example, after I assert (27), you might say:

(29) Yeah, it doesn't have a scratch on it!

and if your assertion is accepted, it may become common ground that we're talking about the token. In this case, I am not committed, by my prior assertion of (27), to the type being beautiful. For example, I can say, without contradicting what I said earlier:

(36) In general the F50 is an ugly model, but this one is beautiful.

Alternatively, you might follow my assertion of (27) by saying:

(37) Yeah, this was the model that got me into Italian sports cars,

and if this assertion is accepted, it may become common ground that we're talking about the type. In this case, I am not committed, by my prior assertion of (27), to the token being beautiful, and I can say, without contradicting what I said earlier:

(38) Too bad this particular one has such an ugly paint job.

King does briefly consider a view with the right shape to explain these facts: that of Chris Barker, who takes the common ground to be a set of world/discourse pairs.⁵⁸ As King notes, Barker's account takes the update for (27) to leave open four possibilities:

- (a) The discourse is about the token, and both the type and the token are beautiful.
- (b) The discourse is about the type, and both the type and the token are beautiful.
- (c) The discourse is about the token, and only the token is beautiful.
- (d) The discourse is about the type, and only the type is beautiful.

⁵⁸ Chris Barker, "Negotiating Taste," *Inquiry* 56, nos. 2–3 (2013): 240–57. This account is formally quite similar to mine. The differences are philosophical: Barker is still attracted to the idea that discourses are parts of worlds (246), and although he expresses reservations about that idea, he lacks a good philosophical account of the distinct roles of the world and discourse parameters of the pairs. That is what I have sought to provide with the expressivist account.

King thinks this is a problem for Barker's view, since "intuitively," only worlds in which both the type and the token are beautiful survive the update.⁵⁹ But that is just what I have been questioning: We can see that there are natural continuations of the conversation that are compatible with both (c) and (d).

⁵⁹ King, *Felicitous Underspecification*, 65.

We have been exploring what account of communication can make sense of vagueness and the more general phenomenon of felicitous underspecification. In chapter 1, we looked at various ways of saving the standard picture, on which a vague assertion is a proposal to add a standard proposition—a condition on states of the world—to the common ground. We concluded that this picture cannot be salvaged. In chapter 2, we considered a modification of the standard picture on which vague assertions put in play a cloud of standard propositions. This idea, too, we dismissed as inadequate. We concluded in chapter 3 that we need to revise our assumptions about assertoric contents. They are not conditions on worlds, but joint constraints on states of the world and plans. The machinery of plan expressivism, we argued, allows us to give an adequate account of vagueness and felicitous underspecification in general.

This whole discussion has taken place at the level of *pragmatics*. We have been talking about the speech acts people make in using sentences like “That test tube is tall.” We have a theory of the *contents* of these assertions (they can be represented as sets of world/hyperplan pairs) and of the *force* of making an assertion (to assert a content is to propose to update the common ground by ruling out world/hyperplan pairs inconsistent with the content). What we don’t have yet is an account of the *meanings* of words like “tall,” or sentences like “That test tube is tall,” that would fit this story. In this chapter, we will look for a semantics that fits our pragmatics.

4.1 The Task of Semantics

Before we get into the weeds, it will be useful to take a step back and think about what is required of a semantic theory. What are the constraints on an account of linguistic meaning, and how does such an account fit into the larger project of explaining communication?

The basic notions we have used in our pragmatic theory—assertion, content, common ground—have no essential connection to language. Speakers can mean things—and presumably make assertions—without using language at all. If I walk into a shop in a foreign country and hold up a battery with an inquisitive look, I will be understood as asking where the batteries are, and the shopkeeper's pointing will be understood as an assertion about their location. Contextual clues and gestures suffice here to make our communicative intentions plain. Communication, then, isn't essentially linguistic. However, most of our communication is linguistically mediated, because by deploying our shared knowledge of language, we can vastly expand the range of meanings we can get across. One might be able to ask where the batteries are by making a gesture, but asking whether the shopkeeper is a recent immigrant is going to be very difficult without language, and asking about the chemical composition of the batteries will be all but impossible.

Language, then, is a tool for expanding the range of communicative intentions we can get across to an audience. But the tool works only when the speaker and the hearer have common knowledge of linguistic meaning.¹ My knowledge of English is of little use in communicating with someone who speaks only Turkish, and my knowledge of the meaning of "ketch" is of little use in communicating with someone who knows what "ketch" means but thinks *I* mean something else by it. Semantics, as a discipline, attempts to describe the content of this common knowledge. What is it that speakers of English know, and take for granted that others know, about the words they use? They may not know anything about pigs, but they know that "pig" is true of pigs, that the plural form "pigs" is used when there is more than one pig, that "large" is an adjective and modifies nouns, that "snort" is a verb in third-person plural present tense, and so on. They also know the meanings of the infinitely many expressions one

¹ James Higginbotham, "Truth and Understanding," *Philosophical Studies* 65, no. 1/2 (1992): 3–16.

can get by combining these basic elements in different ways: for example, “large pig,” “large pigs snort,” and “large pigs snort loudly when you poke them.” A semantic theory for a natural language (or a fragment of one) is an account of the meanings of all of its syntactically well-formed expressions. This specifies the *content* of the semantic knowledge of a speaker who knows the language.

Summing up, then: A semantic theory should tell us things about the basic expressions and modes of combination in a language that can be combined to yield knowledge of the meanings of arbitrary sentences of the language. As for what the meaning of a sentence is, the main constraint we have so far is that it must be something that speakers of the language know (perhaps tacitly), and know in common, and this common knowledge must be useful for figuring out what a speaker means to assert (or suppose, or request, or...) by using that sentence.

If the meanings of sentences just *were* assertoric contents, then it would be easy to see how common knowledge of the meaning of a sentence would be useful for grasping the content a speaker means to assert. However, on most views of assertoric content, this simple picture won't work. The problem is that sentences can have the same assertoric content without having the same role as components of other sentences. For example,

(1) It's snowing

and

(2) It's snowing now

can be used to assert the same content (that it is snowing at the time of utterance), but they do not have the same meaning, because replacing one by the other in a larger sentence can make a difference to assertoric content:

(3) When I awake, I won't know whether it's snowing.

(4) When I awake, I won't know whether it's snowing now.

Here (3) describes future ignorance of a future fact, while (4) describes future ignorance of a present fact. Even if we understand what a sentence does when it is used literally to make an assertion, that might not

be enough to understand how it behaves in the antecedent of a conditional, or embedded under a temporal or modal operator. For this reason, philosophers have distinguished the *compositional semantic value* of a sentence—its meaning in context—from its *assertoric content*.² Pragmatics needs the latter notion, while the job of semantics is to systematize the former.

The notion of compositional semantic value, then, must answer to two independent constraints. On the one hand, it must be rich enough to explain how the meanings of complex expressions depend on the meanings of simpler ones—and this means it may need to encode more distinctions than assertoric content. But it must also be appropriately related to assertoric content. Whatever fine-grained objects play the role of compositional semantic values, we must be able to explain how knowledge of compositional semantic values can help us discern the assertoric content of a speaker's utterance. If, for example, the compositional semantic values were just sequences of natural numbers, it would be mysterious how knowing these could put one in a position to discern a speaker's meaning. There must be a route from compositional semantic value to assertoric content.

In truth-conditional semantics (currently the dominant paradigm), there is a well-recognized way of meeting these two constraints. The compositional semantic value of a sentence is a function from a *context* and an *index* to a truth value.³ An *index* is just a package of parameters that affect the evaluation of expressions and can be shifted by expressions (for example, modal operators shift the world, temporal operators the

² See, for example, Michael Dummett, "Truth," *Proceedings of the Aristotelian Society* 59, no. 1 (1959): 141–62; Michael Dummett, *Frege: Philosophy of Language* (Harvard University Press, 1973); David Lewis, "Index, Context, and Content," in *Philosophy and Grammar*, ed. Stig Kanger and Sven Öhman (Reidel, 1980); Jason Stanley, "Modality and What Is Said," *Language and Mind* 16 (2002): 321–44; Dilip Ninan, "Semantics and the Objects of Assertion," *Linguistics and Philosophy* 33, no. 5 (2010): 335–80; and Brian Rabern, "Against the Identification of Assertoric Content with Compositional Value," *Synthese* 189, no. 1 (2012): 75–96.

³ Paolo Santorio has given some persuasive arguments that this should just be a function from an index to a truth value: "Context-Free Semantics," *Oxford Studies in Philosophy of Language* 1 (2019): 208–39; "Reference and Monstrosity," *Philosophical Review* 121, no. 3 (2012): 359–406. I am happy to embrace this suggestion, but it isn't essential for the point I'm making here.

time, and quantifiers the assignment).⁴ In effect, then, the compositional semantics defines truth at a context and index. How do we get from this to an assertoric content? We do it by defining a notion of *truth at a context*, removing the relativity to an index. For example, if an index includes a world, a time, and a location, then we might say:

- (5) A sentence S is true at a context c iff S is true at $c, \langle w_c, t_c, l_c \rangle$, (where w_c, t_c, l_c are the world, time, and location of the context c).

In earlier work, I have called this step the *postsemantics*, distinguishing it from the compositional semantics proper, which is just concerned with truth at a context and index.⁵ Together, the compositional semantics and the postsemantics are sufficient to define the condition under which a sentence is true at a context. And this is something common knowledge of which would help us grasp the speaker's intended assertoric content. If you've uttered a sentence that we commonly know to be true at the context just in case p holds, then I have reason to take you to have asserted that p .⁶

We can think of the notion of truth at a context, then, as an *interface* between compositional semantics and pragmatics: It connects the technical notion of compositional semantic content, which is used to explain how meanings compose, with the pragmatically relevant notion of assertoric content (see figure 4.1). Knowledge of what it takes for a sentence to be true at a context is something that it is plausible to attribute to ordinary speakers, and something that can explain how language is useful in communication.

⁴ What goes into the index will depend on the expressive resources of the language and how they are being handled theoretically: One might find, for example, standards of precision, tastes, epistemic standards, modal bases, locations, and so on.

⁵ John MacFarlane, "Future Contingents and Relative Truth," *Philosophical Quarterly* 53, no. 212 (2003): 321–36; John MacFarlane, *Assessment Sensitivity: Relative Truth and Its Applications* (Oxford University Press, 2014).

⁶ See Lewis, "Index, Context, and Content," p. 80; David Lewis, "Languages and Language," in *Language, Mind, and Knowledge*, ed. Keith Gunderson, vol. 7 (University of Minnesota Press, 1975), section III. In a Stalnakerian framework, p will be the diagonal proposition (previously discussed in chapter 1, section 1.6). Note that the reason that your choice of sentence gives me for taking you to have asserted that p is defeasible; in cases of nonliteral speech, this reason may be defeated, or one may have a stronger reason for taking you to have asserted something else.




Figure 4.1: Truth at a context as an interface.

In addition, semantic entailment or consequence can be defined in terms of truth at a context: S_1 entails S_2 just in case S_2 is guaranteed to be true at any context at which S_1 is true. This definition allows “It is raining” to entail “It is raining now,” even though the compositional behavior of these sentences is different.⁷ We can explain speakers’ entailment judgments by imputing knowledge of the conditions a context must satisfy for the sentences to be true at it.

It is hard to see how the expressivist approach to vagueness defended in chapter 3 can be accommodated within this standard framework for semantic theory. According to the expressivist approach, the context does not single out a particular threshold for counting as “large” or “tall” (even given a reference class). If that is right, then it does not make sense to ask whether a sentence like “Alice is tall” is true at a context. This suggests that a semantic theory that fits our pragmatics will need to have a fundamentally different shape than orthodox truth-conditional semantic theories.

To make this problem more concrete, let us look at one of the dominant approaches to the semantics of gradable adjectives: the degree theory.⁸

⁷ David Kaplan, “Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals,” in *Themes from Kaplan*, ed. Joseph Almog, John Perry, and Howard Wettstein (Oxford University Press, 1989).

⁸ See Christopher Kennedy, “Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives,” *Linguistics and Philosophy* 30, no. 1 (2007): 1–45; Christopher Kennedy and Louise McNally, “Scale Structure, Degree Modification, and the Semantics of Gradable Predicates,” *Language* 81, no. 2 (2005): 345–81. These works develop a line of thought that goes back to the 1970s: See Ewan Klein, “Comparatives,” in *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, ed. Arnim von Stechow and Dieter Wunderlich (Walter de Gruyter, 1991); Manfred Bierwisch, “The Semantics of Gradation,” in *Dimensional Adjectives*, ed. Manfred Bierwisch and Ewald Lang (Springer, 1989); Arnim von Stechow, “Comparing Semantic Theories of Comparison,” *Journal of Semantics* 3 (1984): 1–77; Lars Hellan, *Towards an Integrated Analysis of Comparatives* (Narr, 1981); Max J. Cresswell, “The Semantics of Degree,” in *Montague Grammar*, ed. Barbara Partee (Academic Press, 2014); Pieter A. M. Seuren,

We will see how degree semantics depends on assumptions that the expressivist approach rejects, and then consider how that theory's valuable compositional insights can be salvaged if we reject those assumptions.

4.2 Degree Semantics for Gradable Adjectives

Gradable adjectives are adjectives that admit of comparative and superlative forms, such as “tall” (“taller,” “tallest”), “flat” (“flatter,” “flattest”), and “intelligent” (“more intelligent,” “most intelligent”). What do we want a semantics for gradable adjectives to explain? It ought to contribute to the explanation of why, when asked to bring the “heavier” bag, Sam brings the 50-pound bag rather than the 35-pound one. But the full explanation of this fact will require physiological and psychological assumptions that fall outside the scope of semantics. After all, if Sam misjudges the weights of the bags, or thinks that the speaker is being sarcastic, he may bring the 35-pound bag. For this reason, semanticists normally focus on data that depend only on a speaker's understanding of language, not on extraneous assumptions. These include facts about the entailment relations the speaker recognizes. For example, speakers accept the following entailments, no matter what ancillary assumptions they make:

- (6) X is larger than $Y \Rightarrow Y$ is not larger than X .
- (7) X is the largest F , Y is an $F \Rightarrow X$ is larger than Y .
- (8) X is large, Y is larger than $X \Rightarrow Y$ is large.
- (9) X is a large F , Y is larger than $X \Rightarrow Y$ is a large F .
- (10) X is large for an F , Y is larger than $X \Rightarrow Y$ is large for an F .

On the other hand, these are not accepted as entailments:

- (11) X is large for an $F \not\Rightarrow X$ is large.
- (12) X is large $\not\Rightarrow X$ is large for an F .
- (13) X is large for an $F \not\Rightarrow X$ is large for a $G F$.
- (14) X is large for a $G F \not\Rightarrow X$ is large for an F .

And then there are cases where we might not be quite sure what to say:

“The Comparative,” in *Generative Grammar in Europe*, ed. Ferenc Kiefer and Nicolas Ruwet (Reidel, 1973).

- (15) X is a large F $\Rightarrow$? X is large.
 (16) X is the largest F $\Rightarrow$? X is a large F.

In addition to this data about entailments, there are data about the *distributions* of words. For example, we can say

- (17) The door is completely closed, or
 (18) The cup is completely full.

but not

- (19) #The building is completely large, or
 (20) #The cup is completely small.⁹

On the other hand, we can say

- (21) The building is very large,

but not

- (22) #The door is very closed.

These facts about distribution and entailment are part of the linguistic knowledge of ordinary speakers, and thus part of the explanatory target of semantic theory.

Perhaps the most obvious way to proceed is to start by giving the meaning of the positive form “large,” and then derive the meaning of “larger than,” “largest,” “large for a,” “large compared to,” “very large,” and other compounds embedding “large” from this. There are theories that proceed this way, but for reasons we won’t go into here, the degree semantics takes a different approach.¹⁰ It supposes that both the comparative “larger” and the positive form “large” are composed from a more basic morpheme which I will write as “*large*.” Thus,

- (23) large = *POS large*,

where *POS* is a null morpheme indicating the positive form, while

⁹ The “#” indicates that the sentence is judged to be marked or unacceptable.

¹⁰ For theories that take the positive form as basic, see Ewan Klein, “A Semantics for Positive and Comparative Adjectives,” *Linguistics and Philosophy* 4 (1980): 1–45; Heather Burnett, *Gradability in Natural Language: Logical and Grammatical Foundations* (Oxford University Press, 2017).

$$(24) \text{ larger} = \text{large} + \text{er}$$

$$(25) \text{ largest} = \text{large} + \text{est}.$$

The semantic value of “*large*” is a *degree function* that maps objects to degrees on an abstract scale:

$$(26) \llbracket \text{large} \rrbracket = \text{largeness-of}.^{11}$$

Larger objects get mapped to higher degrees. The semantic value of the comparative ending “-er,” then, is a function that maps a degree function to the meaning of a regular two-place predicate (a function that takes two objects and returns a truth value):

$$(27) \llbracket \text{er} \rrbracket = \lambda d \lambda x \lambda y (d(x) > d(y)).^{12}$$

Thus,

$$\begin{aligned} (28) \llbracket \text{larger} \rrbracket &= \llbracket \text{er} \rrbracket \llbracket \text{large} \rrbracket \\ &= \lambda x \lambda y (\text{largeness-of}(x) > \text{largeness-of}(y)), \end{aligned}$$

as expected. The semantics of the superlative ending “-est” is a function that takes a degree function as argument and returns a predicate that is true of an object only if it the degree function maps it to a higher degree than any other object in the function’s domain:

$$(29) \llbracket \text{est} \rrbracket = \lambda d \lambda x (\forall y \in \text{Domain}(d) (d(x) > d(y))).$$

What about the positive null morpheme *POS*? It needs to convert a degree function—a function from objects to degrees—into a function from objects to truth values. It does this by comparing the degree of the object to a threshold. The threshold is generated by applying a contextually supplied threshold function s_c to the degree function itself. So, for example, $s_c(\text{largeness-of})$ is the minimum degree of largeness an object needs to have for “large” to be true of it. Thus, the semantic value of the positive morpheme *POS* is

$$(30) \llbracket \text{POS} \rrbracket = \lambda d \lambda x (d(x) \geq s_c(d)),$$

¹¹ Here, $\llbracket \cdot \rrbracket$ is a function from expressions to their semantic values.

¹² Read this as: “the function that takes a degree function and two objects and returns the value True if the degree function maps the first object to a higher degree than the second object, and False otherwise.”

and the semantic value of “large” is, then,

$$(31) \llbracket \text{large} \rrbracket = \llbracket POS \rrbracket \llbracket \text{large} \rrbracket \\ = \lambda x (\text{largeness-of}(x) \geq s_c(\text{largeness-of}))$$

Modifiers like “for an ant” or “compared to Sarah” are understood as functions from degree functions to degree functions. They modify a degree function by restricting its domain. So, for example, the semantic value of “large for an ant” is just *largeness-of*_{ants} (the *largeness-of* function with its domain restricted to ants). More generally,

$$(32) \llbracket \text{for an } F \rrbracket = \lambda d (d|_{\llbracket F \rrbracket}).$$

An explicit comparison like “compared to Sarah” is just a special case of this: It restricts the domain of the degree function to just two objects:

$$(33) \llbracket \text{compared to } a \rrbracket = \lambda d \lambda x (d|_{\{x, \llbracket a \rrbracket\}}(x)).$$

Kennedy emphasizes that comparison classes are often implicit—we say “She is tall” when we mean “She is tall for a third-grader”—and that they can be different from the common noun used with the adjective—an expensive BMW need not be expensive *for a BMW*.

This theory can explain entailments (6–10) and nonentailments (11–14). Just to take two examples, (10) is an entailment because if X exceeds the threshold defined for the *largeness-of* function restricted to the Fs, and Y has a higher degree of largeness than X does, then Y must also exceed this threshold. And (11) is not an entailment because “large” and “large for an F” have different degree functions as their semantic values. Even though the latter is just the former restricted to a smaller domain, it is still a different function, and the contextually supplied s_c function can map it to a different threshold (either higher or lower).

Kennedy goes on to explain the distributional data (15–22) by appealing to differences in the *scale structures* of *relative* gradable adjectives, like “large,” which have no natural end point, and *absolute* gradable adjectives, like “full,” which have a natural end point (in this case, being completely full). We won’t get into the details here, but it is important to recognize the scope of phenomena explained by the degree semantics. It’s a good theory; it explains a lot, and we should be loath to give it up.

However, it looks as if we *must* give it up, for two reasons. First, the semantics, and the common knowledge it attributes to speakers,

presupposes that the context of a speech act determines a function s_c that maps every gradable adjective meaning to a threshold. As we have argued in chapters 1–3, that is not a plausible assumption. Speakers do not think that context fixes a particular number of millimeters at which “tall” first applies, and they certainly don’t take this to be common ground. They may think that the context rules out certain thresholds: If it has been agreed that Feynman counts as “tall,” then we’ve ruled out thresholds above his height. So the context of the conversation might confine thresholds to a certain range. But speakers do not take there to be a hidden fact about where the threshold lies within that range. They do not have an attitude of *uncertainty* about the threshold. They regard it as a matter for decision, not discovery (see section 3.1 in chapter 3).

The issue is not one of metaphysics. An argument that, contrary to what speakers believe, the facts about the concrete speech situation do determine an unknown threshold would not suffice to vindicate the use of the s_c function in the semantics. The semantic theory’s account of truth at a context is supposed to match speakers’ understanding of the conditions that a context must satisfy to be true at a context (section 4.1). So unless speakers presuppose that there is a contextually determined threshold for counting as “large,” the degree semantics is not giving a correct account of their understanding.

Nor is this something that has an easy, obvious fix. We needed s_c to give an account of the positive form (the semantics for *POS*), and without s_c , it is not clear how we could do that. But suppose we find a way to patch the degree semantics so it does not presuppose a contextually supplied threshold function s_c . (Perhaps we supervaluate, quantifying over the thresholds in a contextually determined range.) So long as the theory retains the standard shape of a truth-conditional semantics, yielding a definition of truth at a context, it is unclear how it can fit our pragmatics. The standard bridge from truth at a context to assertoric content goes something like this:

1. The speaker is being sincere and literal.
2. So she must mean to assert that things are the way they need to be for the sentence to be true at the context at which it is uttered.
3. It is common knowledge between us that the sentence is true at the context just in case p .

4. So, the speaker means to assert that p .

But this reasoning yields a standard proposition, a condition on worlds, not the kind of assertoric content we need for our expressivist pragmatics, a joint condition on worlds and plans.

It is starting to look as if our expressivist pragmatics will need a radical new semantic theory—one that does not culminate in a definition of truth at a context and does not assume the existence of a contextually determined threshold-setting function. And that is what I concluded in my first explorations of this topic.¹³ There, I proposed a semantics that puts a *hyperplan* parameter in the index, even if nothing shifts this parameter and there is no compositional motivation for it.

However, I now think that there is a better approach. We can keep the degree semantics, almost exactly as Kennedy presents it. With a little bit of reinterpretation, we can make it compatible with our pragmatics and avoid reinventing the wheel.

4.3 Semantics Without Truth at a Context

Here is a way to preserve the core of existing semantic theories while disentangling them from their problematic commitments. As we have seen, we cannot make sense of these theories if we interpret the “ c ” in expressions like “ s_c ” or “true at c ” as referring to a context—a possible speech situation. But we can simply *reinterpret* the notation, taking “ c ” to denote what I will call a “hypercontext” rather than a context. A *hypercontext* is an abstract entity that can provide values to parameters or variables, as contexts were assumed to do, but without the assumption that they correspond to possible speech situations. The name change signals a resolution not to assume that our hypercontexts play the role usually assigned to contexts in pragmatics. In particular, we do not assume that there is a unique hypercontext that is relevant to a speech situation or conversation.

When we think of our postsemantics as defining truth at a hypercontext, rather than truth at a context, we need to rethink the “bridge” between the postsemantics and assertoric content. Given an orthodox

¹³ John MacFarlane, “Vagueness as Indecision,” *Aristotelian Society Supplementary Volume* 90, no. 1 (2016): 255–83.

semantic theory that defines truth at a context, we can tell a fairly simple story about how hearers can use their knowledge of the semantics of a sentence S to determine what a speaker is asserting by uttering S .¹⁴ The hearer can reason as follows: The speaker is presenting S as true in the context and takes it to be common knowledge that S is true at a context just in case that context has such-and-such features, so the speaker must be intending to convey that the situation they are in has such-and-such features (for example, that it will rain the day after, that the speaker is tired at the time of the utterance, or that uranium is radioactive). But this reasoning presupposes that a context, as used in the postsemantics, determines a concrete possible speech situation. Hypercontexts do not correspond to concrete possible speech situations. So we need a new story about how truth at a hypercontext determines an assertoric content, conceived now as a set of world/hyperplan pairs.

A world/hyperplan pair will determine a hypercontext: The world will fix the actual world of the hypercontext, while the hyperplan, together with the world, will fix values for the other contextual parameters, including the formerly problematic s_c . Even if there is no fact of the matter what the threshold for *tall* is, it makes sense to *plan* to set a threshold at a certain point, and since a hyperplan is a maximally determinate plan, a hyperplan will settle all questions about thresholds, either unconditionally or conditionally on the state of the world.¹⁵ We can then formulate a new bridge principle relating hypercontexts and assertoric content:

¹⁴ Lewis, "Index, Context, and Content"; Robert Stalnaker, "Assertion," in *Syntax and Semantics, Vol. 9: Pragmatics*, ed. P. Cole (Academic Press, 1978).

¹⁵ For the latter case, consider a plan to set the threshold for *large apple* to the median diameter of a certain reference class. The world will determine what the median is, so the plan and the world will jointly determine a threshold.

Here, I am ignoring the possibility, floated in chapter 3, section 3.3, that a hyperplan might be "indifferent" between two or more settings of a threshold. If we adopted that idea, we would need some changes in the compositional semantics itself, as s_c would now pick out a set of thresholds rather than a single one. The semantics for *POS* would then need to change to allow a third truth value: "neither." "*POS large*" would be true of an object x if $\delta =_{\text{def}} \text{largeness-of}(x) \geq$ all the thresholds in the set, false if $\delta <$ all the thresholds in the set, and neither if δ is less than some thresholds in the set and greater than or equal to others.

One might wonder what plans have to do with other standard parameters of the context, like the time or agent of the context. We will come back to this issue in section 4.6.

Hypercontexts and assertoric content

The assertoric content associated with a sentence used literally is the set of world/hyperplan pairs that determine some hypercontext at which that sentence is true.

If we adopt this principle and reinterpret the “*c*”s in our formal semantics as denoting hypercontexts, we can keep everything in the formal compositional semantics the same. In particular, we can keep the clause of *POS*, with its essential use of s_c . Saying that a hypercontext determines s_c is compatible with denying that there is a fact of the matter about the value of s_c in a concrete speech situation.

A terminological digression. Some semanticists may say:

The thing you call a *hypercontext*—a sequence of parameter values that has plays certain role in the semantic theory—is just what we call a *context* and distinguish from a concrete speech situation. Indeed, we allow that it may be indeterminate, in a given concrete speech situation, which context is the context of use.

We have already seen this way of talking in von Fintel and Gillies (discussed in chapter 2, section 2.1).¹⁶ It can also be found in Alex Silk, who says that it can be indeterminate which abstract context represents the discourse context;¹⁷ in Christopher Kennedy and Malte Willer, who say that “some aspects of (formal) context that matter for truth and falsity are *essentially* underdetermined, in the sense that their underdetermination persists *no matter* the particulars of the utterance situation”;¹⁸ and in Fabrizio Cariani and Paolo Santorio, who say that while there is always a world of the context, “it may be indeterminate what context the utterance takes place in, and hence which world is supplied to the semantics.”¹⁹

¹⁶ Kai von Fintel and Anthony S. Gillies, “‘Might’ Made Right,” in *Epistemic Modality*, ed. Andy Egan and Brian Weatherson (Oxford University Press, 2011).

¹⁷ Alex Silk, *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics* (Oxford University Press, 2016), 45.

¹⁸ Christopher Kennedy and Malte Willer, “Familiarity Inferences, Subjective Attitudes and Counterfactual Contingency: Towards a Pragmatic Theory of Subjective Meaning,” *Linguistics and Philosophy* 45, no. 6 (2022): 1411.

¹⁹ Fabrizio Cariani and Paolo Santorio, “Will Done Better: Selection Semantics, Future Credence, and Indeterminacy,” *Mind* 127, no. 505 (2018): 142.

In talking with such semanticists, we can use the following translation manual: our “hypercontext” = their “context,” and our “context” = their “concrete speech situation.” This terminological difference should not affect any of the substantive issues in play. If we don’t think that a context is determined by a concrete speech situation, then we need to explain how hearers can use their knowledge of what it takes for a sentence to be true at a context to discern the speaker’s intended assertoric content. If, as I have suggested, we think of indeterminacy as a matter of being undecided by our linguistic and conceptual plans, then we can solve this problem by thinking of the assertoric content as the set of world/hyperplan pairs that determine some context at which the sentence is true. Those who don’t theorize indeterminacy as indecision will have to give a different story, but some story is owed.

4.4 Have Meanings Gotten Too Thin?

The meanings described by this theory are *thin*. Knowing the meaning of “tall,” on this view, requires knowing that it is a measure of height, that it is the positive form related to the comparative “taller,” that it can be modified by “very” but not by “completely.” But beyond structural facts of this kind, nothing more substantive need be known.

For example, someone could know everything about “tall” that is encoded in the semantic theory, and still think it makes sense to classify *all* American males as “tall for an American male.” Nothing about the *semantics* forbids planning to put the threshold at the very bottom of the range, or leaving it open that one could decide to do this. In this way our theory differs sharply from Delia Graff Fara’s, which requires that people in the extension of “tall (for an American male)” be significantly taller than the “typical” height for this reference class.²⁰ “Tall” is more flexible than such constraints allow, and we can easily find cases where one might want to apply “tall” even to things that are below average height, and below the typical height, for the reference class. Suppose, for example, that there is one building that is 300 feet tall, eight that are 50 feet tall, and four that are 10 feet tall. Only the tallest building is significantly taller than the average height (about 57 feet) or the typical height (50

²⁰ Delia Graff Fara, “Shifting Sands: An Interest-Relative Theory of Vagueness,” *Philosophical Topics* 28, no. 1 (2000): 45–81, discussed previously in chapter 1, section 1.4.

feet). Still, it seems to me that there are at least two reasonable ways to use “tall” here—one on which it applies only to the 300-foot building, and the other in which it also applies to the 50-foot buildings.

One might concede this but still want to build in *some* constraints: For example, doesn’t the meaning of “tall” at least rule out classing 5-foot-tall men as “tall” (for an adult American male)? But I don’t think there is a strong motivation for building in further, minimal constraints of this kind. To be sure, a Martian who knew the meaning of “tall” but had no other experience of terrestrial life might have trouble aligning its use of “tall” with the expectations of human speakers. But that’s because fluent conversation depends on more than just knowledge of linguistic meaning. As John Searle famously noted, someone who asks for a hamburger in a restaurant will expect to get one that is a certain size (say, 4–10 inches in diameter, and wider than it is thick) and not encased in a Lucite cube.²¹ Our Martian’s ignorance of this fact may engender a failure of understanding, but we don’t think this is a reason to suppose that “hamburger” really means “ground beef patty of a certain size, and not encased in a Lucite cube.” A meter-wide hamburger *is* a hamburger, and so is a Lucite-encased one.

Similarly, in ordinary human conversations, it will be commonly accepted that 5-foot-tall men are not tall (for an adult American male).²² Lacking this common knowledge, a speaker who knows the linguistic meaning of “tall” may still fail to understand and be understood. As meanings get thinner, this kind of nonlinguistic conversational background knowledge plays an increasingly important role in explaining communication.

Geoffrey Nunberg has argued that the derogatory aspects of slurs are not part of their conventional meaning, any more than the fact that the word “comely” is old-fashioned is part of its meaning:

Clearly the fact that *redskin* is derogatory is an arbitrary fact about that word that speakers have to know in order to use or understand it. Somebody who is ignorant of that fact is deficient in her knowledge of English, not in her knowledge of racism, American Indians or the rules of conversational interaction. But there are a lot of things we know about

²¹ John Searle, “Literal Meaning,” *Erkenntnis* 13, no. 1 (1978): 207–24.

²² As observed in chapter 3, section 3.5, this can be common knowledge even if it is not commonly known what the greatest height is that counts as tall.

words which affect the way we use them but which are all the same not themselves part of its meaning.²³

I want to say something similar about people who don't realize that an American man classified as "tall" is probably more than 5 feet tall. Speakers' use of these words is guided not just by their knowledge of linguistic meaning, but also by their knowledge of how the words have typically been used in different contexts.

4.5 Beyond Threshold Indeterminacy

We have focused so far on the meanings of gradable adjectives, but vagueness, and the more general phenomenon of felicitous underspecification, extends far beyond gradable adjectives. Jeffrey King's main examples of felicitous underspecification—for example, "those guys" or "local"—do seem to involve indeterminacy about the setting of a contextual (or, we might now say, hypercontextual) parameter.²⁴ We can handle them the same way we are handling indeterminacy in gradable adjectives—namely, by saying that the context is compatible with a range of hypercontexts. But not all cases of vagueness or indeterminacy can be explained this way. For example, common nouns like "couch" and "chair" seem vague. At any rate, one can imagine a sequence of small changes connecting a paradigm couch and a paradigm love seat, or a paradigm chair and a paradigm bar stool. Their vagueness does not seem to involve indeterminacy about the setting of a threshold on a degree scale (degrees of "couchness"?). Nor is there any other obvious contextual parameter in play here. Similarly, any indeterminacy in verbs like "run" or "walk" seems to have its source in the basic lexical items themselves. Other examples are not hard to find: "planet,"²⁵ "hexagonal,"²⁶ "momentum,"²⁷

²³ Geoffrey Nunberg, "The Social Life of Slurs," in *New Work on Speech Acts* (Oxford University Press, 2017), 264.

²⁴ See section 1.5 in chapter 1.

²⁵ Mike Brown, *How I Killed Pluto: And Why It Had It Coming* (Spiegel & Grau, 2010).

²⁶ David Lewis, "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8, no. 3 (1979): 352.

²⁷ Hilary Putnam, *Representation and Reality* (MIT Press, 1988), 9–10.

“mass,”²⁸ “athlete,”²⁹ “cake,”³⁰ and “human.”³¹ Presumably just about every nonfunctional word will make the list.

Indeed, indeterminacy about threshold-setting doesn't exhaust the indeterminacy even of gradable adjectives. Here is an example, adopted for my purposes from Charles Travis.³² Luisa is planning her wedding. It's an autumn wedding, and she says she wants *red leaves* on the dinner table with the flowers. Since the leaves on the trees have not yet started to change color, the florist takes some green leaves and paints them red. Has Luisa gotten what she asked for? That depends. There are different “understandings” (to use Travis's term) of “red leaf.” On some understandings, the leaf must be naturally red; on others, it is the surface color that counts, even if it is artificial, so one can have a *green leaf that is painted red*. Normally, when we talk about a red leaf, we mean one that is naturally red. But one *can* certainly use “red leaf” to mean one that is red on the surface. There is a degree of flexibility about what a speaker can express using “red leaf,” and about what update to the common ground a speaker can be proposing in asserting “Those are red leaves.”

Recall that on the Kennedy-style degree semantics for *red*, the semantic value of “red” (for a leaf) is a partial function mapping leaves to degrees of redness. So the flexibility at issue here would be about which such function is associated with “red” (for a leaf). This is flexibility about the semantic value of a lexical item, not about the setting of a contextual parameter or variable. To put the point another way, Travis's sort of indeterminacy affects not just the positive form “red” but the comparative “redder.”³³

²⁸ Hartry Field, “Theory Change and the Indeterminacy of Reference,” *Journal of Philosophy* 70 (1973): 463–67; Mark Wilson, “Predicate Meets Property,” *Philosophical Review* 91, no. 3 (1982): 564.

²⁹ Peter Ludlow, *Living Words: Meaning, Underdetermination and the Dynamic Lexicon* (Oxford University Press, 2014), 78.

³⁰ David Edmonds, “Cake or Biscuit: Why Jaffa Cakes Excite Philosophers,” BBC News, February 20, 2017, <http://www.bbc.com/news/magazine-38985820>.

³¹ Colin Barras, “We Don't Know Which Species Should Be Classed as ‘Human’,” BBC News, February 11, 2017, <http://www.bbc.com/earth/story/20160111-what-is-it-that-makes-you-a-human-and-not-something-else>.

³² Charles Travis, “Pragmatics,” in *A Companion to the Philosophy of Language*, ed. Bob Hale and Crispin Wright (Blackwell, 1997).

³³ Cf. Timothy Williamson, *Vagueness* (Routledge, 1994), 155–56.

We see similar flexibility with “tall” when we consider whether, for example, a 927-foot building with no spire on top is tall to a higher degree than a 925-foot building with a 111-foot spire on top. (In 1929, the answer to this question would determine whether the Bank of Manhattan or the Chrysler Building was the world’s highest building.) For different purposes, we might use a mapping of buildings to degrees of height that includes spire height in the ranking, or one that does not. Similarly, in evaluating “large” (for a city), we might base degrees of largeness on either population or physical extent.³⁴ How should we handle this flexibility in our semantic and pragmatic theories?

Travis seems to assume that in any given context, there is a particular “understanding” that captures the speaker’s meaning. Grasping the update the speaker is proposing is a matter of recognizing which understanding is intended. However, the assumption that a speaker and hearer can coordinate on a particular understanding of lexical items like “tall” may often be just as implausible as the assumption that they can coordinate on particular settings for contextual parameters (such as a threshold or reference class). In many cases, the speaker’s intention is not sufficiently determinate to select among the candidate understandings. In asking for “red leaves,” Luisa may not have thought about whether painted-red leaves would be sufficient. In describing Los Angeles as a “large city,” one may not have decided whether largeness should be measured by population or physical extent. And in asking which is the “tallest building,” one may not have thought about whether spires and antennas should be included in height measurements. In these cases, it seems to me, there is no single intended understanding; instead, the speaker is undecided among several candidate understandings. And even in cases where speakers have a single preferred understanding in mind, it will often be impossible for them to expect hearers to recognize which one it is. In such cases, they will be unable to mean something involving this preferred understanding, because speaker meaning essentially involves an intention that one’s communicative intentions be recognized.

The issues here closely resemble the issues we discussed with contextual completions, and at the level of pragmatics, the same solution seems to work in both cases. We can model indeterminacy in the semantic

³⁴ Kennedy, “Vagueness and Grammar.”

value of a lexical item as a kind of practical indecision. The mental state expressed by an utterance of

(34) The Chrysler Building is taller than the Bank of Manhattan,

then, can be a hybrid practical-doxastic state, and the associated update to the common ground can consist in the ruling out of certain combinations of worlds and hyperplans. In a context where it is common knowledge how tall each building is, both with and without spires, an assertion of (34) may serve merely to express a plan to count spires (or at least this particular spire) in calculating total height. In a context where it is common knowledge that spires will count, but not how tall the buildings are, (34) will express an ordinary belief that the Chrysler Building (with its spire) is taller than the Bank of Manhattan. And in a context where there is both uncertainty about the heights and indecision about whether to count spires, (34) will express a hybrid state, ruling out, for example, deciding not to count the spire *while* believing that the spireless height of the Chrysler Building is less than the spireless height of the Bank of Manhattan.

In the pragmatics, then, we can handle lexical underspecification in just the same way we handle contextual underspecification. But what should we say in the semantics? As we have seen, existing semantic theories contain a parameter that determines the threshold, and with some reconceptualization, we can see this as being constrained by speakers' plans for threshold-setting rather than determined by the context. But, as Michael Glanzberg emphasizes, these theories do not have a parameter for the scale itself, which they typically just specify homophonically ("degrees of tallness" or perhaps "degrees of linear extent"), if they specify it at all.³⁵ Compare (26), which picks out a degree function as *largeness-of*. Homophonic clauses are also used to give the semantics of nouns like "couch":

(35) $\llbracket \text{couch} \rrbracket = \lambda x.x \text{ is a couch}$.³⁶

³⁵ Michael Glanzberg, "Explanation and Partiality in Semantic Theory," in *Metasemantics: New Essays on the Foundations of Meaning*, ed. Alexei Burgess and Brett Sherman (Oxford University Press, 2014).

³⁶ Cf. Irene Heim and Angelika Kratzer, *Semantics in Generative Grammar* (Blackwell, 1998), 64.

Glanzberg suggests that we think of these homophonic parts of semantic theories as pointers to concepts, which are themselves outside the scope of semantic theorizing. The semantics of gradable adjectives can explain why “large” has an antonym “small”; why “completely full” is acceptable but “completely large” is not; why “A is large” and “B is larger than A” together imply “B is large”; and why “It is not closed” implies “It is open,” while “It is not large” does not imply “it is small.” But when it comes to the differences in meaning between “large,” “smart,” “long,” and “fuzzy,” semantic theories just have nothing to say. Glanzberg supposes that there *is* a story about these differences, but this story is part of psychology rather than semantics. The homophonic clauses are thus a kind of pointer to something outside of semantics proper.

I think something like this is right, but if the general view I’ve been pushing is correct, then we can’t adopt Glanzberg’s tolerant attitude to the homophonic bits of semantic theories. If what we count as a couch, or as a greater degree of largeness, is subject to planning indeterminacy, then words like “couch” and “largeness” should not be used in the semantic metalanguage, which tries to describe how the language works from a neutral standpoint. (For similar reasons, we should not explain the meaning of “I” in our semantic theory by saying: “‘I’ refers to me.”) So, homophonic semantic clauses like (26) or (35) must be replaced. But what can we replace them with?

One natural idea would be to add a parameter of the hypercontext corresponding to each of the homophonic bits. For example, instead of

$$\llbracket \text{large} \rrbracket = \text{largeness-of},$$

we’ll write

$$\llbracket \text{large} \rrbracket^c = c_{\text{largeness-of}}.$$

In the former, “largeness-of” is being used with its ordinary meaning; in the latter, it just serves to label a particular parameter of the hypercontext, whose value can be constrained by the evolving plans in a conversation. Similarly, instead of

$$\llbracket \text{couch} \rrbracket = \lambda x.x \text{ is a couch},$$

we’ll have

$$\llbracket \text{couch} \rrbracket^c = c_{\text{couch}}.$$

In this way, we’ll convert the homophonic elements in the compositional clauses for lexical items to hypercontextual parameters.

The meanings described by this semantic theory are *really* thin. They encode structural features of the sort that play a real explanatory role in semantic theories, but their “substance”—what distinguishes “tall” from “smart”—is drained out of the semantic theory itself, and pushed out to the background nonsemantic knowledge. It’s not that “tall” and “smart” have the *same* meaning on this account, because their semantic clauses point to different parameters of the hypercontext, representing different shared concepts. But the detailed knowledge about how these concepts differ—about what factors matter for arranging things on the *tallness* scale or the *smartness* scale, respectively—is now counted as extrasemantic knowledge.

Is this going too far? In section 4.1 we said that semantics tries to say what any two speakers of English know in common, giving as an example that “pig” is true of pigs. But given what we’ve said so far, someone could count as knowing the semantic meaning of “pig” without realizing that “pig” is true of pigs. Suppose they routinely apply “pig” to dogs and “dog” to pigs. We might be inclined to say that they don’t know what “pig” means. They might still be inclined to assent to the sentence “‘Pig’ is true of pigs,” but given how they use the word “pig,” we might doubt that this sentence expresses knowledge that “pig” is true of pigs. “‘Pig’ is true of pigs,” they might say—adding, “There goes one now,” as a dog runs by.

Certainly, such a person suffers from a deficiency that will greatly hamper communication. But is it a deficiency in their grasp of the meaning of “pig”? Or, rather, is it a deficiency in their knowledge about what pigs are like? I am inclined to say, following Burge, that if we can describe them as participating in a group that intends to use a common concept, a group that links their claims and thoughts under the label “pig,” then they are using “pig” with the same meaning as anyone else.³⁷

4.6 Panvariabilism

At this point, we have a view on which the semantic value of pronouns, common nouns, and gradable adjectives are all relative to a hypercontext.

³⁷ Tyler Burge, “Individualism and the Mental,” *Midwest Studies in Philosophy* 4, no. 1 (1979): 73–122. In a footnote, Burge declines to affirm even that a man who thinks that “orangutan” applies to a fruit drink “cannot correctly and literally be attributed a belief that an orangutan is a kind of fruit drink” (91).

And it's not just that the hypercontext influences their semantic values: The hypercontext *fully determines* their semantic values:

$$\llbracket couch \rrbracket^c = c_{couch}$$

$$\llbracket large \rrbracket^c = c_{largeness-of}$$

$$\llbracket they_i \rrbracket^c = c_{they_i}$$

A hypercontext, then, is a function that maps expressions to values. That is essentially what a variable assignment does. The only difference is that a variable assignment only assigns values to *variables*. But what is a variable, anyway? Semantically, variables are just linguistic elements whose values are completely given by an assignment. So it is tempting to do some theoretical unification and reconceive things more radically. If we think of all these expressions as variables, then we don't need hypercontexts; we just need assignments.

If we think of things this way, we'll represent

(36) That's cheap for a Mercedes

as

(37) $POS_{t_1}(for-a(n_1)(g_1))(x_1)$

Here, "cheap" is represented by the variable g_1 (ranging over degree functions); it differs from "fast," "smart," and "tall" only in that they are *different* variables ranging over degree functions. Similarly, "Mercedes" is represented by the variable n_1 (ranging over semantic values of common nouns); it differs from "cow," "car," and "bicycle" only in being a different variable. "That" is represented by the variable x_1 (ranging over individuals); it differs from "I," "you," "he," and "she" only in being a different variable. Finally, t_1 , a variable ranging over functions from degree functions to thresholds (degrees), does the work of Kennedy's contextually determined s_C .³⁸

³⁸ We think of this variable as being attached to the null morpheme *POS*. That is a significant difference from standard degree semantics, which has a global, contextually determined threshold-setting function that is used in the semantics of every occurrence of *POS*. This might look like an advantage for the standard approach, if we can't hear a true reading of, for example, "He's tall for a basketball player but not tall for a basketball player." What prevents us from giving different values to the threshold-

Accepting (37) puts a joint constraint on worlds and assignment functions, ruling out world/assignment pairs such that (37) is not true at that world and assignment. If we think of hyperplans as, *inter alia*, plans for assigning values to variables, then accepting (37) imposes a joint constraint on worlds and hyperplans. So this conception of the semantics fits well with our pragmatics.

This approach erases the distinction between cases where the indeterminacy affects a contextual parameter (for example, gradable adjectives or plural pronouns) and cases of lexical indeterminacy. Both are now just cases of indeterminacy about assignments of values to variables. This is a good thing. Our hypercontext semantics handled only contextual underspecification, leaving us with no account of lexical underspecification. The panvariabilist semantics, by contrast, promises a uniform treatment of both.

This advantage notwithstanding, one might reasonably think that treating *every* lexical item as a variable is just crazy. Surely, it is part of the meaning of “I” that it denotes the speaker (or, more generally, the “agent of the context”). If the semantics is charged with recording facts about meaning, then this should be reflected somewhere in the semantics. But if “I” functions as a variable, then the only thing the semantics says about it is that its denotation is the value that the assignment function assigns to that variable. Similarly, it would seem that the differences between “he” and “she,” and between “car” and “cow,” should be recognized in the semantics. These words have distinctively different meanings, which we can’t capture just by taking them to be distinct variables!

In responding to this objection, we can build on a strategy pioneered by early advocates of variabilism, who advocated it only for pronouns and proper names.³⁹ We can concede that understanding the differences

setting functions governing the two positive occurrences of “tall,” if they have distinct variables? To prevent this, we need to say that every occurrence of *POS* comes with the same free variable. That may seem *ad hoc*, but as we will see in what follows, that is just how we will deal with uses of the pronoun “I” by the same agent, following Santorio, “Reference and Monstrosity.”

³⁹ Sam Cumming, “Variabilism,” *Philosophical Review* 117, no. 4 (2008): 525–54; Santorio, “Reference and Monstrosity”; Dilip Ninan, “Counterfactual Attitudes and Multi-Centered Worlds,” *Semantics and Pragmatics* 5 (2012): 1–57. Their motivations are discussed later in this chapter (section 4.8).

between the pronouns “I,” “he,” and “she” is part of our linguistic competence. Someone who doesn’t know that “I” refers to the speaker, or that “he” and “she” constrain the gender of the referent, is failing to grasp something essential to the meanings of these words. But, as Paolo Santorio observes, this knowledge doesn’t need to enter into the compositional semantics. Instead, it can be deployed at the *postsemantic* level. Recall that *postsemantics* is the part of a linguistic theory that connects the notion of truth at an index, defined by the compositional semantics, with contexts. In a traditional Kaplan- or Lewis-style semantic theory, for example, the postsemantics might say that a sentence is true at a context if it is true at the index of the context. Santorio proposes that the rule stating that the reference of “I” is the agent of the context can be part of the postsemantics. On his view, the context determines a distinguished variable assignment, the “assignment of the context,” and we can stipulate that if the variable α is used to represent a use of “I” by a speaker, the assignment of the context assigns that speaker to α . So, in the compositional clauses, we just have variables, whose values are given by the assignment component of the index; but in the postsemantics, there are constraints on what can count as “the assignment of the context.”⁴⁰

This reply was sufficient for the original, restricted variabilists, who just claimed that proper names and pronouns should be treated as variables. But we cannot adopt it without revision because it presupposes something we are committed to rejecting: that there is an “assignment of the context.” Because we want to handle underspecification, we can’t assume that the context determines a particular assignment. Indeed, there is nothing in our picture that corresponds to the postsemantics, because we no longer think it makes sense to define truth at a context. Nonetheless, we can hold on to the idea that understanding the meanings of these pronouns requires grasping certain *constraints* on variable assignments—for example, that the variables representing uses of “I” should get mapped to the respective speakers, or that variables representing uses of “she” should be mapped to females—which are imposed *outside of* the semantic theory’s compositional clauses. These constraints become part of the evolving common ground of a conversation. When “I”

⁴⁰ Santorio, “Reference and Monstrosity,” 385.

is used, for example, it becomes common ground that the assignment assigns the speaker to the variable representing “I.”⁴¹

Similarly, although “car” and “cow” are represented by free variables of the same type, the common ground will include different constraints on assignments to these variables. For example, the assignment to “cow” might be constrained to be a subset of the assignment to “animal.”⁴² Plausibly, this constraint is one we might think that anyone who grasps the meaning of “cow” must accept. But other constraints on the assignment won’t be “analytic” in this sense. For example, if we have been using two pronouns, “he₁” and “he₂,” and we learn that he₁ and he₂ are one and the same, then the common ground will constrain the assignment to map “he₁” and “he₂” to the same object. That constraint represents a bit of synthetic knowledge. Other cases may be harder to classify. For example, in a particular conversation, it might be common ground that the assignment to “athlete” must be a subset of the assignment to “human.” It may have been decided in the conversation that no animal, regardless of its speed or agility, can count as an athlete.⁴³ I do not think it would be plausible to say that anyone who thinks that a horse can be an athlete just doesn’t grasp the standing meaning of “athlete.” The *Oxford English Dictionary* says that “athlete” means “a competitor in any sporting event that requires physical skill, strength, endurance, etc.” That seems to leave things open, until we decide whether a horse race is a sporting event and the horse a competitor. This is just the kind of case that makes one’s grip on the notion of analyticity start to falter. However, for the purpose of understanding how communication works when there is felicitous underspecification, it is of no importance whatsoever to divide the constraints the common ground imposes on the assignment into the categories “analytic” and “synthetic.”

In general, these constraints will not uniquely determine the denotations of variables. Perhaps in the case of “car” and “cow,” the constraints

⁴¹ For a similar view about demonstratives, see Ethan Nowak, “No Context, No Content, No Problem,” *Mind & Language* 36, no. 2 (2021): 189–220.

⁴² Or, more properly, a “subintension.” The assignment will map predicate variables to intensions: functions from worlds to extensions. *X* is a subintension of *Y* just in case at any world *w*, the extension of *X* at *w* is a subset of the extension of *Y* at *w*.

⁴³ Ludlow, *Living Words*, 78, describes hearing an argument on talk radio about whether a racehorse could be an athlete.

will be fairly stringent, while for “athlete” or “cookie” or “planet,” they will be looser and embody more indecision. What constraints are established will depend on the details of the conversation and on the background knowledge the parties have of each other. For example, it may be common ground in a discussion between two biologists that “red” should be applied to tree leaves based on the underlying color, regardless of paint, whereas among artists, it may be common ground that what counts for the application of “red” is what kind of light the surface reflects.

4.7 Assignment-Sensitive Contents

We now have a semantic theory that treats all lexical items and contextual parameters as variables in an appropriate semantic category. The “context change potential” of a sentence—the way a literal assertion of it will affect the common ground, if accepted—is determined by the set of world/assignment pairs that satisfy it.

Indeed, it is convenient for some purposes to think of the common ground as a set of world/assignment pairs instead of a set of world/hyperplan pairs. A set of world/hyperplan pairs will determine a set of world/assignment pairs (including just those pairs that are compatible with some world and hyperplan in the set), so we can still explain what it is for a given set of world/assignment pairs to be the common ground by appealing to shared Gibbardian hybrid doxastic-planning states. But now we get a more straightforward connection between the semantics and the pragmatics. We can think of the content of a sentence as the set of world/assignment pairs that satisfies it. The essential effect of an assertion is simply to intersect its content with the common ground.

Some may object that this way of thinking about assertoric content individuates contents in the wrong way. Intuitively, if I say

(38) That₁ is cute,

pointing to a penguin, I have not asserted the same thing as I would have if I’d been pointing at a café. But on the account we have been exploring, both claims can be represented as

(39) X_1y_1 ,

and both have the same content—the set of world/assignment pairs (w, g) at which the object g assigns to y_1 belongs to the set of objects it assigns to X_1 at w . (Of course, if I demonstrate the penguin *and* the café in the same conversation, *different* variables must be used, but nothing hangs on whether “ y_1 ” or “ y_2 ” is used when the penguin is demonstrated.)

It is important to distinguish two questions here. The first is whether our account can make sense of ordinary certain intuitive same-content judgments. It has no trouble explaining the judgment that if I say, “That is cute” (demonstrating a café) and then later, “That is cute” (demonstrating a penguin), I have made claims with different contents. In this case, different variables would be used. The account only has trouble with *counterfactual* judgments of content identity, for example: “The content asserted by (38) would have been the same whether a penguin or a café was demonstrated.” Here, I would like to say that accounting for such intuitive same-content judgments is not an important theoretical goal. As David Lewis observes, we use different criteria of same-saying in different cases, and I think the same goes for same-asserting and same-believing.⁴⁴ Sally says, “I am hungry.” Manuel says, “I am hungry.” Did they both assert the same thing? For some purposes, we might say yes; for other purposes, no. On Monday Sally said, “It will surely be sunny tomorrow.” Did she assert the same thing she asserted on Sunday when she said, “It will surely be sunny tomorrow”? For some purposes we might say yes, for other purposes, no. We might have local purposes where it matters how we individuate contents, but these vary depending on our purposes. There is no particular reason to believe that the way of individuating contents that is optimal for a systematic theory of meaning and communication will match the ways we commonly employ in everyday life.⁴⁵

⁴⁴ Lewis, “Index, Context, and Content,” 97.

⁴⁵ Cumming, who endorses variabilism but not *panvariabilism*, seems willing to think of contents as assignment-sensitive: “I prefer to think of open propositions as an intermediate layer of content. This way I can have the character (or linguistic meaning) of a sentence, relative to a context, determine (or constrain) the use to which its context-sensitive expressions are put, and thus the open proposition it expresses in that context. The open proposition may then be evaluated for truth at the assignment and world provided by the context, just as formerly a closed proposition was evaluated for truth at the world of the context” (Cumming, “Variabilism,” 542; for an elaboration of how an “open proposition” can be the content of a belief, see p. 545). This is essentially the

The second question is whether our account, which assigns a common content to all assertions of (38), even when different objects are demonstrated, can explain the obvious differences in the pragmatic significance of these assertions. This question is serious, but our account has the resources to respond to it. The pragmatic significance of an assertion depends not just on its content, but also on the common ground. We have seen this already in connection with Chris Barker's example

(40) Feynman is tall.⁴⁶

If it is common ground that Feynman is, say, 192 cm tall, then the practical effect of asserting (40) is to constrain the threshold for counting as tall such that everybody whose height is at least 192 cm counts as tall. If it is not common ground how tall Feynman is, but it is common ground that people at least 192 cm count as tall, then the practical effect of asserting (40) is to convey the information that Feynman is at least 192 cm tall. And if neither of these things is common ground, but the common ground is both uncertain about Feynman's height and undecided about the threshold, then asserting (40) has a hybrid effect, constraining *combinations* of possibilities as to Feynman's height and plans about where to place the threshold. On our account, the *content* asserted by an utterance of (40) is the same in all three scenarios; it is the differences in the common ground that account for these differences in practical significance.

In the same way, we can explain why asserting (38) has different significance in different contexts, even though the same content is asserted. In the two scenarios we have described, the common ground will include different constraints on the assignments to y_1 and X_1 .

From a formal point of view, the picture we have been developing looks similar to some versions of dynamic semantics. In particular, Irene Heim proposed a view on which the context is represented as a set of world/assignment pairs, plus a set of distinguished variable indices indicating

picture I am recommending, except that I don't think there is an "assignment and world provided by the context," and hence no such thing as a "closed content."

⁴⁶ Chris Barker, "The Dynamics of Vagueness," *Linguistics and Philosophy* 25, no. 1 (2002): 1–36, discussed previously in chapter 3, section 3.2.

which variables are “free.”⁴⁷ The compositional clauses will then associate a similar set of world/assignment pairs (and a set of distinguished indices) with each sentence. Have we simply reinvented an old idea?

No. These formal similarities are superficial. Although both our approach and Heim’s represent a conversational context as a set of world/assignment pairs, they give very different stories about the significance of these sets. In Heim’s semantics, we existentially quantify over assignments at the level of the whole discourse to obtain an ordinary truth condition. The whole discourse is true in world w if there is an assignment a of values to the variables such that the pair (w, a) belongs to the common ground’s set of world/assignment pairs. Thus, for example, the discourse

(41) A man came in. He was wearing a coat. He sat down next to a woman.

x_1 is a man. x_1 came in. x_1 was wearing a coat. x_2 is a woman. x_1 sat down next to x_2 .

is true in w just in case there is an assignment of values to the variables x_1 and x_2 such that, in w , the value of x_1 is a man wearing a coat, the value of x_2 is a woman, and x_1 sits next to x_2 . Here, we certainly would not want to say that (41) expresses a mental state involving indecision about which people “a man” and “a woman” denote. Semantically, these indefinites involves quantification; it’s just that on Heim’s picture, the quantification is applied at the discourse level rather than the sentence level.

By contrast, the free variables associated with lexical items in our picture are never existentially closed. When the common ground does not constrain a variable to a unique value, we have semantic indecision, which gets spelled out as indecision in plans.

So, although Heim’s approach and ours both involve representing common grounds as constraining variable assignments at a world, they differ greatly in the significance they give to these constraints. On Heim’s approach, the constraints are part of a theory that systematizes truth

⁴⁷ For an elementary exposition, see Seth Yalcin, “Dynamic Semantics,” in *Routledge Companion to Philosophy of Language*, ed. Gillian Russell and Delia Graff Fara (Routledge, 2012), section 3.

conditions at the level of whole discourses. On the expressivist approach, the constraints represent a kind of practical indecision, and we never pass from the constraints to truth conditions, even at the whole-discourse level.

One might want to combine a Heimian approach to indefinites with an expressivist approach to contextual underspecification. In that case, one would have to distinguish two types of variables. The variables used in indefinites would then be subject to existential closure, while the variables used to represent lexical items and contextual parameters would remain merely constrained. We will not pursue these complications here.

4.8 Epistemic Shift

I have motivated panvariabilism as a way to handle semantic underdetermination. This was not the motivation of the original variabilists, who were happy to suppose that context provided a unique variable assignment. The original motivation for variabilism was to allow the denotations of pronouns to be “shifted” by attitude verbs, conditionals, and epistemic modals. I now want to suggest that, if these considerations carry any weight, they support not just a restricted variabilism (about proper names and pronouns), but *panvariabilism*.

Consider John Perry’s example of the amnesiac scholar who has forgotten who he is.⁴⁸ He knows he is either Lingens, who haunts the Main Library at Stanford, or Lauben, who hangs out in Widener Library at Harvard. Looking around, he sees the stacks of a library, but he can’t tell which library it is. He might give voice to his uncertainty by uttering:

- (42) I might be Lingens and I might be Lauben.
- (43) If this is the Main Library, I’m Lingens.
- (44) If this is Widener Library, I’m Lauben.

If we make the orthodox assumption that “I” refers to the agent of the context, then it is hard to see how these sentences could all be true. On standard assumptions about modals and conditionals, the truth of (42) requires that “I am Lingens” be true relative to one possible world and

⁴⁸ “Frege on Demonstratives,” *Philosophical Review* 86, no. 4 (1977): 492.

“I am Lauben” relative to another. But on David Kaplan’s account, “I” is directly referential; its referent depends only on the context of use, not on the world of evaluation.⁴⁹

(45) $\llbracket I \rrbracket^{c,w,t,g} = \text{the agent of } c.$

So shifting the world should not change the referent of “I.” Similarly, Saul Kripke taught us that the referent of a proper name does not depend on the world of evaluation. So shifting the world should not change the referent of “Lauben” or “Lingens.” Given these assumptions, one of “I am Lingens” and “I am Lauben” is true at all possible worlds, and the other is true at no possible worlds. And this makes it hard to understand how (42), (43), and (44) can all be true.

One natural solution would be to say that epistemic modals and indicative conditionals shift the context, so that the consequents of (43) and (44) are evaluated relative to different contexts, and (42) in effect quantifies over contexts compatible with some body of information. This would allow us to keep the Kaplanian semantics for “I,” at the cost of rejecting Kaplan’s claim that English does not contain “monsters”—constructions that cause their constituents to be evaluated relative to a shifted context.⁵⁰ However, this would only be a temporary victory. Similar examples of epistemic shift can be constructed for proper names, which on Kaplan’s semantics are not context-sensitive. Imagine an ancient astronomer saying:

(46) Hesperus might be Phosphorus, but it might not be Phosphorus.

⁴⁹ One of Kaplan’s important points is that “I” is not synonymous with “the speaker,” whose denotation does depend on the world of evaluation.

⁵⁰ Kaplan, “Demonstratives,” 511. In the past decades, we have learned that there are quite a few languages with context-shifting expressions, including Zazaki, Nez Perce, and Korean; for a thorough treatment, see Amy Rose Deal, *A Theory of Indexical Shift: Meaning, Grammar, and Crosslinguistic Variation* (MIT Press, 2020), which describes the patterns of context-shifting we find in these languages and develops a systematic theory to explain them. For example, in Nez Perce, one can say something like “Last week, John said I will go to town tomorrow” and mean that last week, John said that *he* would go to town tomorrow (Deal, *A Theory of Indexical Shift*, 3). But English does not exhibit the kinds of systematic indexical shift that these languages do. So there is something to Kaplan’s idea that English, at least, contains no “monsters,” and we should try not to erase this difference between English and languages like Nez Perce.

- (47) If Pythagoras's calculations are correct, then Hesperus is Phosphorus.
- (48) If Anaxagoras's calculations are correct, then Hesperus is not Phosphorus.

Again, it seems that we can only make these sentences true if the denotations of "Phosphorus" and "Hesperus" can be shifted by the epistemic modal and conditional antecedent. But shifting the context will not effect such a shift because these terms are not context-sensitive.⁵¹

We could reconsider Kaplan's ban on monsters in English *and* his view that proper names are context-sensitive.⁵² But there is an alternative—one that does not require changes to Kaplan's basic semantic assumptions. The alternative is to suppose that pronouns and proper names function as *variables*, and the epistemic constructions shift the assignment in conjunction with the world.⁵³ This allows the denotation of "I" to differ between the two scenarios envisaged in our example while keeping its semantics directly referential. And in general, it offers a solution to Frege's puzzle that does not, like Frege's own solution, require that names and pronouns denote something different in epistemic contexts than they do outside them.⁵⁴

I don't want to argue here that this is the best solution to Frege's puzzle. I simply want to point out that *if* this application motivates variabilism for pronouns and proper names, it ought to motivate variabilism for *every lexical category*. For example, suppose that Jane knows the words "ketch" and "yaw" but has forgotten which one is the one with the

⁵¹ At any rate, they are not context-sensitive on Kaplan's account: Kaplan, "Demonstratives," p. 562; David Kaplan, "Afterthoughts," in *Themes from Kaplan*, ed. Joseph Almog, John Perry, and Howard Wettstein (Oxford University Press, 1989), p. 598.

⁵² Some alternative theories do take proper names to involve a context-sensitive element. Tyler Burge, "Reference and Proper Names," *Journal of Philosophy* 70, no. 14 (1973): 425–39, for example, takes proper names to be predicates; when they are used as singular terms, there is an implicit demonstrative, whose referent depends on the context.

⁵³ See Cumming, "Variabilism"; Santorio, "Reference and Monstrosity"; and Ninan, "Counterfactual Attitudes and Multi-Centered Worlds."

⁵⁴ Frege took names and pronouns in the complements of attitude verbs to denote their customary senses, rather than their customary referents (e.g., people). Gottlob Frege, *Collected Papers on Mathematics, Logic, and Philosophy*, ed. Brian McGuinness (Basil Blackwell, 1984), 166.

mizzenmast forward of the rudder post. Looking at a two-masted yacht in the marina, *Zephyr*, she notices that its mizzenmast is forward of the rudder post, and thinks:

- (49) *Zephyr* might be a ketch, but it might not be a ketch.
- (50) If a ketch has the mizzenmast forward of the rudder post, then it's a ketch.
- (51) If a ketch has the mizzenmast aft of the rudder post, then it's not a ketch, but a yawl.

If (49) is true, there should be epistemically possible worlds w_1 and w_2 such that *Zephyr* falls into the extension of "ketch" at w_1 , but not at w_2 . But for this to be so, "ketch" must have a different intension in the first conjunct of (49) than it does in the second conjunct. In the first conjunct, it must have an intension that applies to (at least some) yachts with the mizzenmast forward of the rudder post, while in the second conjunct, it must have an intension that applies to no such yachts. Jane is *certain* that the boat she sees has its mizzenmast forward of the rudder post, so there are no epistemically possible worlds where it does not. We cannot explain the truth of (49), then, just by talking about the shifting of worlds. We need the epistemic modals to shift, not just the world of evaluation, but the intension of "ketch." It can do that if we take "ketch" to be a variable and allow the modal to shift the assignment, just as we did with the pronouns in (42) and the proper names in (46).

It is easy to find similar examples for other lexical categories, such as verbs:

- (52) She might have kowtowed to the emperor, and she might not have.
- (53) If kowtowing requires actually touching your forehead to the ground, then she did not kowtow.
- (54) If kowtowing requires only lowering your forehead near the ground, then she did kowtow.

The reasons for embracing variabilism as a solution to the problem of "epistemic shift," then, are reasons for accepting *panvariabilism*.

4.9 Semantic Shift

Panvariabilism also gives us a natural way to make sense of a phenomenon theorized by Alexander Kocurek, Ethan Jerzak, and Rachel

Rudolph (KJR) as “convention shifting.”⁵⁵ KJR argue that we can hear two distinct readings of sentences like

- (55) If Pluto were a planet, there would be at least a dozen planets in the solar system.
- (56) Pluto could have been a planet, but it isn't.
- (57) Pluto used to be a planet, but since 2006, it hasn't been one.⁵⁶

On the orthodox readings (drummed into us by Kripke and Kaplan and legions of their loyal followers), (55) asks us to consider a scenario in which Pluto is a planet, in the sense accepted today—that is, a scenario in which Pluto clears its orbital neighborhood and has enough gravity to force it into a spherical shape. On this reading, (55) says that if Pluto had been able to clear its orbital neighborhood, there would have been several additional planets. Similarly, (56) says that there is a possible past in which Pluto cleared its orbital neighborhood. And (57) is understood as saying that Pluto used to clear its orbital neighborhood, but doesn't anymore. On these readings, none of the sentences is plausible. But KJR argue that we can *also* get readings of these sentences in which the modal or temporal expressions shift the conventions for the use of the word “planet.” So, we can read (55) as saying that conventions that would count Pluto as a planet would also count several other bodies, currently not counted as planets, as planets. We can read (56) as saying that there are some possible conventions that would have counted Pluto as a planet. And we can read (57) as saying that the conventions in place before 2006 did count Pluto as a planet, but the conventions in place since 2006 do not. To get these readings, they argue, we need to suppose that counterfactual conditionals and modal and temporal constructions can shift conventions, not just worlds.

I want to concede that we can hear the alternative readings of (55–57), and we need some account of them. The question is whether these readings motivate an index parameter for hyperplans or hyperconventions,

⁵⁵ Alexander W. Kocurek, Ethan Jerzak, and Rachel Etta Rudolph, “Against Conventional Wisdom,” *Philosophers' Imprint* 20, no. 22 (2020): 1–27.

⁵⁶ In 2006, the International Astronomical Union (IAU) adopted a definition of “planet” that requires planets to have cleared their orbital neighborhoods. Pluto does not meet this condition. It meets the other conditions, but so do a number of bodies that were not traditionally considered planets, such as Ceres, Haumea, and Makemake. For an engaging account of the discoveries leading up to the IAU's decision and the reasoning behind it, see Brown, *How I Killed Pluto*.

as KJR suppose. It is not obvious to me that they do, for these examples seem relevantly similar to a number of others that don't seem to involve any kind of convention shifting, like the following:

(58) If this year had not been a leap year, today would be March 1.

(59) Today is February 29, but last year it was March 1.

(60) If I'd been Einstein, I would have aced my physics quiz.

I find these all ambiguous: On the standard reading, where "this year," "today," and "I" rigidly designate a particular year, day, and person, these are defective in some way (either false or possibly unevaluable or vacuously true, depending on one's approach to counterpossibles). But on another reading, they express substantive truths. These readings do not involve shifted conventions for using words. We might think they shift the context at which the indexicals are evaluated. But since variables have now taken over all the work formerly done by contexts (or hypercontexts) in the compositional semantics, we can say instead that they shift the assignment together with the world. We've already seen precedent for this in Santorio's account of epistemic shift and related work by Cumming and Ninan.⁵⁷

Once we have introduced assignment-shifting, a panvariabilist doesn't need any further machinery to explain the cases KJR explain by appealing to "convention shifting." In (55), for example, the counterfactual can shift the intension of "planet," now conceived as a variable, by shifting the assignment. We need not posit a parameter of the index representing linguistic conventions.

We can use the same approach to handle attitude ascriptions. KJR note that

(61) Alpha thinks that Pluto is a planet

doesn't imply that Alpha thinks that Pluto clears its orbital neighborhood.⁵⁸ It is difficult to explain this if we treat "thinks" as a modal operator quantifying over doxastic alternatives. KJR suggest that "thinks" shifts conventions as well as doxastic alternatives. But although this will help with (61), it does not help with very similar cases, like

⁵⁷ Santorio, "Reference and Monstrosity"; Ninan, "Counterfactual Attitudes and Multi-Centered Worlds"; Cumming, "Variabilism."

⁵⁸ Kocurek et al., "Against Conventional Wisdom," 6.

- (62) Alpha thinks he might be in the Main Library and might be in Widener Library,⁵⁹ or
 (63) Alpha thinks this one is that one.⁶⁰

Allowing attitude verbs to shift the assignment handles these cases well. If we're panvariabilists, it handles (61) in exactly the same way.⁶¹

Other devices that seem to shift conventions include scare quotes and locutions like "literally," "metaphorically," or "in *x*'s sense".⁶²

- (64) The "freedom fighters" will arrive at the city today.
 (65) The ice queen is hot, but not in the literal sense.
 (66) In Quine's sense of "planet," there are nine planets.

A panvariabilist can handle all of these as assignment-shifting constructions.

4.10 Adieu Diagonalization

In section 1.6 we asked whether Stalnakerian diagonalization could solve the problems posed by felicitous underspecification, and we answered this question in the negative. We can now ask the converse question: Now that we have a general mechanism for handling felicitous underspecification, do we still need diagonalization?

I believe the answer is no. All the cases diagonalization was designed to help with can be seen by the panvariabilist as cases in which the context does not fully constrain the assignment. Let's reconsider Stalnaker's example: Speaker and Hearer both hear a woman talking in the next room, and Speaker says:

⁵⁹ Santorio, "Reference and Monstrosity."

⁶⁰ Cf. Cumming, "Variabilism."

⁶¹ Compare Brian Rabern, "Semantic Monsters," in *Routledge Handbook of Linguistic Reference*, ed. Heimir Geirsson and Stephen Biggs (Routledge, 2021), 525: "There are already treatments of attitude verbs that treat them as operators that shift the assignment of values to variables (e.g. Cumming 2008, Ninan 2018, Santorio 2012, Rabern forthcoming). Shifting the assignment is tantamount to shifting the 'interpretation' of the variables, and thus assignment shifters are a close, though perhaps more innocent, cousin to meaning-shifting devices. Given that these proposals only work for singular terms, a natural generalisation would be to have epistemic contexts shift the interpretation of general terms such as 'eye doctor' and 'optometrist' as well."

⁶² Kocurek et al., "Against Conventional Wisdom," 8.

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

On Stalnaker's view, the proposition expressed by (67) at the context of use is either the necessarily true proposition (if the woman in the next room is in fact Gabor or Anscombe) or the necessarily false proposition (if she is someone else). Neither of these can be the content that Speaker is asserting, given the plausible pragmatic principle that it must be common ground between Speaker and Hearer what proposition is being asserted.⁶³ Hence Speaker and Hearer need to coordinate on another proposition as the asserted content. Because they have a shared conception of which proposition would be expressed by (67) in each possible world, they are in a position to coordinate on the diagonal proposition: The proposition that is true at a world w just in case the proposition that would be expressed by (67) were w actual is true at w .

This case does not involve semantic indecision. Speaker has, after all, decided to use "that" to refer to the woman talking in the next room, whoever she may be, and Hearer recognizes this. However, if we take the contents of assertions to be sets of world/assignment pairs, we can handle the Anscombe/Gabor case without resorting to diagonalization. On our new approach, (67) expresses a set of world/assignment pairs that contains a pair (w, g) just in case the object g assigns to the variable associated with "that" in (67) is Gabor or Anscombe in world w . We can simplify this by supposing there is just one variable, the one associated with "that," and three objects, G (Gabor), A (Anscombe), and C (Tricia Nixon Cox).⁶⁴ There are then three possible assignments, g_G (which assigns Gabor to the variable), g_A (which assigns Anscombe), and g_C (which assigns Cox). We can also suppose there are just three possible worlds, w_G (in which Gabor is speaking), w_A (in which Anscombe is speaking), and w_C (in which Cox is speaking). Then the content of (67) is

$$\{(w_G, g_G), (w_A, g_G), (w_C, g_G), (w_G, g_A), (w_A, g_A), (w_C, g_A)\},$$

⁶³ Robert Stalnaker, *Context and Content: Essays on Intentionality in Speech and Thought* (Oxford University Press, 1999), 88, puts this principle slightly differently: "The same proposition is expressed relative to each possible world in the context set." Because the context set contains the worlds left open by the common ground, this amounts to the assumption that it is common ground which proposition is being asserted.

⁶⁴ On our full view, the proper names will also be treated as variables, but that doesn't matter for our purposes.

for (67) is true so long as “that” gets assigned either Anscombe or Gabor, no matter what the state of the world. To assess the effect of asserting this content, we need to think about what the common ground looks like after it is asserted, but before Hearer has decided whether to accept it.⁶⁵ We can assume that Speaker intends to use “that” to refer to the person speaking in the next room, and that Hearer recognizes this intention. So it is common ground that the referent of “that” is the speaker in the next room—that is,

$$\{(w_G, g_G), (w_A, g_A), (w_C, g_C)\}.$$

If Hearer accepts the assertion, this common ground is updating by intersection with the content of (67), yielding

$$\{(w_G, g_G), (w_A, g_A)\}.$$

The effect of the assertion is thus to remove one possible combination of assignment and world: the assignment of Tricia Nixon Cox to “that” and the world where she is the speaker. We get a prediction similar to Stalnaker’s, but without resorting to diagonalization; instead, we update with the literal content of the assertion.

Making contents constraints on world/assignment pairs, rather than just worlds, therefore gives us the tools to solve two rather different kinds of problems: problems of *undecided* reference (felicitous underspecification) and problems of *decided, but uncertain* reference.

4.11 Conclusion

Our task in this chapter was to say how semantic theories need to change to accommodate the sort of pragmatics argued for in chapters 1–3. Standard semantic theories for gradable adjectives assume that the context picks out a threshold—an assumption we rejected. We can get around this problem by reconceiving the “c” parameter in our semantics as a “hypercontext.” But a similar problem arises for lexical indeterminacy, and it cannot be solved in the same way. As a general solution, we have proposed treating *all* lexical items as variables, whose semantic values are given by the assignment. The resulting view, *panvariabilism*, requires

⁶⁵ See Stalnaker, “Assertion,” 86.

no new semantic machinery, and it can be motivated independently as a way of dealing with epistemic and semantic shift.

However, adopting panvariabilism means conceding that much of the common stock of assumptions about word usage that allows us to communicate effectively is up for negotiation in the context of a conversation, and outside the scope of the knowledge the semantic theories aims to capture. In this respect, it resembles Richard Montague's model-theoretic semantics. In an influential critique, Ernie Lepore complains that Montague semantics tells us only that "'Barbara sekoilee' is true in Finnish if and only if whatever 'Barbara' picks out is one of the things 'sekoilee' is true of"—a fact someone could know without having any idea what they would be saying if they uttered "Barbara sekoilee."⁶⁶ A Davidsonian semantics, by contrast, would tell us the sorts of things we need to know in order to interpret utterances of the sentence: that "Barbara" denotes Barbara, that "sekoilee" is true of something just in case it is confused, and that the whole sentence is true just in case Barbara is confused. What Lepore says about Montague's view could just as well be said about panvariabilism, which treats "Barbara" and "sekoilee" as variables and not only declines to specify the variable assignment of the context, but denies that there is such a thing.

To this objection, I believe Richmond Thomason gives the proper response (in his introduction to Montague's collected papers):

A central goal of this account is to explain how different kinds of meanings attach to different syntactic categories; another is to explain how the meanings of phrases depend on those of their components. ... *But we should not expect a semantic theory to furnish an account of how any two expressions belonging to the same syntactic category differ in meaning.* "Walk" and "run," for instance, and "unicorn" and "zebra" certainly do differ in meaning, and we require a dictionary of English to tell us how. But the making of a dictionary demands considerable knowledge of the world.⁶⁷

A dictionary collects lore about the expected common ground insofar as it regards the usage of words. Aside from categorial information, all of this

⁶⁶ Ernest Lepore, "What Model Theoretic Semantics Cannot Do?" *Synthese* 54, no. 2 (1983): 177.

⁶⁷ Richmond Thomason, "Introduction," in *Formal Philosophy: Selected Papers of Richard Montague*, ed. Richmond Thomason (Yale University Press, 1974), 48, emphasis in the original.

is defeasible, and it often happens that a common ground will depart from these standard expectations. None of it is in the proper domain of semantic theory.

It may be illuminating to compare the panvariabilist account with the theory Donald Davidson sketches in "A Nice Derangement of Epitaphs."⁶⁸ Davidson is interested in cases where word usage diverges locally from standard expectations: for example, cases of unintentional malapropism (Archie Bunker saying "We need a few laughs to break up the monogamy") or simple usage mistakes (calling someone by the wrong name or confusing "ketch" and "yaw"). He proposes that in such cases, we need to distinguish two semantic theories, one describing the semantic expectations speakers bring to a conversation (the *prior theory*) and one describing the common expectations they converge on dynamically within the conversation (the *passing theory*). The panvariabilist, similarly, distinguishes between the common ground of assumptions about usage before the conversation and the one that governs the conversation as it proceeds. One might wonder whether it makes a difference which way we describe what happens: Is it all the same whether we say that there are two semantic theories, or one semantic theory and two different sets of common ground assumptions?

I believe it does make a difference, and here is why. A semantic theory in Davidson's style needs to yield truth conditions, and it must therefore specify a condition for an object to fall into the extension of a predicate. However, at a given point in a conversation, we may have only *constraints*. For example, we may have decided that *Gs* definitely count as *Fs* and *Hs* do not, but we may be leaving it open (and subject to further refinement) whether *Ks* count as *Fs*. This is easy to account for on a panvariabilist account: Just as the participants' shared view of the facts can be partial and evolve as the conversation proceeds, so the constraints they accept on word usage can be partial and evolve. But it is hard to see how to account for this on a Davidsonian account: The passing theory has to state truth conditions and can't leave things open. Davidson limited his attention to cases in which the participants mean something definite but nonstandard—for example, cases where speakers

⁶⁸ Donald Davidson, "A Nice Derangement of Epitaphs," in *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, ed. Ernest Lepore (Blackwell, 1986).

mean *yawl* by “ketch.” But in a very common kind of case, speakers do not mean anything definite; they are not *sure* what “ketch” means, or they want to remain neutral about what it takes for something to count as a “planet.” Panvariabilism offers a framework in which such cases are perfectly ordinary, not something embarrassing to be explained away or swept under the rug.

According to the view presented in chapter 3, indeterminacy in our thought and talk is just practical indecision—indecision about how to use our words and concepts. In asserting

(1) The Bank of Manhattan is taller than the Chrysler Building,

one is proposing that we constrain both our linguistic usage and our view of the world in such a way that (1) is accepted. Depending on what is already common ground, such an assertion might amount to a purely descriptive claim (if it is already accepted that the height of a building should be measured to the top of its highest architectural spire), or a proposal about how to use “taller” (if the relevant heights and architectural details are already common ground), or a hybrid proposal that cannot be cleanly factored into descriptive and plan-laden parts.

The idea that assertions often aim to push plans for using words has tended to be discussed in recent literature under the heading of “metalinguistic negotiation.” Those who claim that metalinguistic negotiation is pervasive in discourse have the right basic picture, but they need a better frame to hang it in. The right frame for their insights, I will urge, is the plan-expressivist pragmatics I have been developing in this book. Lacking this frame, theorists of metalinguistic negotiation have been forced to accept some untenable conclusions.

In this chapter, I will focus on the most influential analysis of metalinguistic negotiation, due to David Plunkett and Timothy Sundell.¹ Plunkett and Sundell defend three main claims:

¹ David Plunkett and Timothy Sundell, “Disagreement and the Semantics of Normative and Evaluative Terms,” *Philosopher’s Imprint* 13 (2013): 1–37; David Plunkett and

- In a wide variety of cases that look like canonical disagreements, the locus of disagreement is not the asserted propositions (which are compatible), but normative metalinguistic attitudes expressed in making the assertions.
- Hence the *argument from disagreement*—which passes from a premise about the existence of disagreement to a conclusion that the disagreeing parties “mean the same thing” by a key term—is invalid.
- Standard semantic and pragmatic theories already have the resources to explain what is going on in cases of metalinguistic negotiation.

I will argue that all three claims are false. Accommodating the general phenomenon of metalinguistic negotiation requires changing standard semantic and pragmatic theories in the way suggested in chapters 3 and 4. Once we have done that, we will see that the locus of disagreement in the cases of interest is the asserted proposition, that the opposing parties are *not* asserting compatible propositions, and that the argument from disagreement can be vindicated.

At the end of the chapter, I will briefly consider three alternative accounts of metalinguistic negotiation: Keith DeRose’s single-scoreboard semantics, Alex Silk’s discourse contextualism, and Justin Khoo’s quasi-indexicalism. Although I think all three accounts are improvements on Plunkett and Sundell’s, I will argue that they all have shortcomings that are addressed by the expressivist semantics and pragmatics developed here.

5.1 Plunkett and Sundell’s Account

In explaining the basic idea of Plunkett and Sundell’s account, I’ll start with what I think are the most favorable cases for motivating their view. The first concerns the word “chips,” which is used differently in American and British English:

Timothy Sundell, “Metalinguistic Negotiation and Speaker Error,” *Inquiry* 64 (2021): 142–67; David Plunkett and Timothy Sundell, “Metalinguistic Negotiation and Matters of Language: A Response to Cappelen,” *Inquiry* 68, no. 10 (2021): 3123–47; Timothy Sundell, “The Tasty, the Bold, and the Beautiful,” *Inquiry* 59, no. 6 (2016): 793–818.

(2) American: The burgers come with chips.

Brit: No, they don't come with chips. They come with *crisps*.²

The second comes from a talk radio discussion heard by Peter Ludlow:³

(3) Host: Surely Secretariat, who won all three Triple Crown races, was one of the greatest athletes of the last century.

Caller: No, Secretariat is not one of the greatest athletes of the last century. Only humans can be athletes.

The third involves threshold negotiation for a gradable adjective, "rich":

(4) Kamala: People who make \$400,000 a year are rich.

Donald: No, they're not rich. They're upper middle class. Rich people have yachts and fly on private planes.

These examples all express some kind of disagreement, which is made explicit by the word "No" in the replies. Applying the argument from disagreement to these examples will yield the conclusions that American and Brit mean the same thing by "chips" in (2), that Host and Caller mean the same thing by "athlete" in (3), and that Kamala and Donald mean the same thing by "rich" in (4). Plunkett and Sundell reject these conclusions. In (2), they say, the two speakers obviously mean different things by "chips," so they are not really contradicting one another. Similarly, in (3) the speakers mean different things by "athlete," and in (4) they mean different things by "rich."⁴ Because the speakers in these cases use terms in different senses, they aren't actually asserting incompatible

² Plunkett and Sundell, "Disagreement and the Semantics of Normative and Evaluative Terms."

³ Peter Ludlow, *Living Words: Meaning, Underdetermination and the Dynamic Lexicon* (Oxford University Press, 2014). The subject of the conversation was a *Sports Illustrated* list of the best athletes of the twentieth century that included Secretariat, the great American racehorse.

⁴ A context-sensitive word like "local" or "today" has the same *standing meaning* even when it is used in relation to different places or days. On most accounts, gradable adjectives like "rich" are context-sensitive expressions of this sort (see chapter 4, section 4.2). Here, "mean the same thing" should be understood as "mean the same thing, together with its contextual supplementation." In this sense, people who use "that" to demonstrate different objects do not mean the same thing, even though both are using the word with its standard meaning.

propositions. "Many genuine disagreements are expressed *via* disputes in which the speakers literally express *compatible* contents."⁵

If the two parties are asserting compatible contents, why do we perceive them as disagreeing? Here, Plunkett and Sundell observe, quite correctly, that disagreeing with someone does not require claiming that what they said is false. Consider these examples:

(5) Anne: Your wife is very beautiful.

Bjorn: No, we're not married.

Here, "No" targets a presupposition of Anne's assertion rather than its main content.

(6) Alfred: Either Wilson or Heath will be the next prime minister.

Bernard: I disagree, it will be either Wilson or Thorpe.⁶

The two disjunctions asserted by Alfred and Bernard are compatible; what is at issue is whether Heath or Thorpe is a serious contender.

(7) Ani: She is good-looking.

Beth: No, she is *gorgeous*.

Here, Beth's claim doesn't contradict Ani's; presumably gorgeous people are good-looking. Beth is rejecting Ani's claim as too weak.

(8) Amy: It will rain tomorrow.

Bob: No, there's only a 50 percent chance.

Bob is not committed to the falsity of Amy's claim; he is only objecting to its assertability.

(9) Abe: The Sioux used to live on this land.

Beth: No, the *Lakota* used to live here.

Beth prefers the Lakota's own name for themselves to the name "Sioux," which derives from their enemies' derogatory term for them ("Nadowes-sioux," meaning "little snakes" in Ojibwa). But she doesn't object to the content of the claim.

⁵ Plunkett and Sundell, "Disagreement and the Semantics of Normative and Evaluative Terms," 12, emphasis in the original.

⁶ H. P. Grice, *Studies in the Way of Words* (Harvard University Press, 1989), 64–65.

(10) Arthur: He had sex with her without her consent.

Ben: No, he *raped* her.

Here, the two claims are basically synonymous; Ben is objecting to the overly clinical manner in which Arthur is describing the situation.

In each of these cases, the locus of disagreement is something other than the truth of the first speaker's assertion. Clearly, then, we cannot infer from the use of "No" that the second speaker rejects the content of the first speaker's assertion: "No" can be used to reject assertions for many reasons besides the falsity of their contents. Still, one might think the combination of "No" and the negated form of the original sentence is harder to dismiss. Examples (2–4) all have that form, while none of the examples I have just given do. So we do need a specific story about what the locus of disagreement is in these cases.

According to Plunkett and Sundell, the disagreement in (2–4) concerns a normative, metalinguistic issue: what we ought to mean by "chips" (or "athlete," or "rich"). Note that the issue is not what we *do* mean by these words; by assumption, the two speakers mean *different* things by them. But in using them this way, Plunkett and Sundell suggest, they are expressing a view about how these words *ought* to be used. Brit in (2) is advocating for a usage of "chips" on which it denotes thick, chewy fried potatoes, not thin, crispy fried potatoes. Caller in (3) is advocating for a usage of "athlete" that excludes nonhuman animals. And Donald in (4) is advocating for a very high standard for who counts as "rich."

To summarize, then, Plunkett and Sundell claim that in cases like (2–4), which they call "metalinguistic negotiations," the following two conditions obtain:

Assertoric Compatibility

The two speakers are asserting contents that are not contradictory.

Expressive Incompatibility

In doing so, they are expressing normative, metalinguistic attitudes that *are* contradictory.

The *locus of disagreement* in these cases—the thing the two parties disagree about—is an unexpressed normative metalinguistic claim, not an asserted content.

For example, in (3),

each speaker literally expresses a true proposition given the concept they in fact express with their term. But beyond that, the speakers pragmatically advocate for the concept that they are using and in virtue of which they assert those propositions. Thus, their metalinguistic dispute reflects a genuine disagreement about how to use the word “athlete.”⁷

It follows that the argument from disagreement is invalid: “Philosophers are wrong to reason from the fact that a particular exchange involves a genuine disagreement to the thesis that the speakers involved in that exchange mean the same things by the words they use.”⁸

5.2 How Metalinguistic Disagreements Can Be Substantive

Some philosophers have objected that taking disagreements to be covertly metalinguistic trivializes them. And to be sure, the disagreement in (2) *does* seem like a paradigm example of a trivial disagreement—one we’d characterize as “merely semantic.” But Plunkett and Sundell mean their analysis to extend to a whole range of disagreements that don’t seem trivial at all, such as the following:

(11) Barack: Waterboarding is torture.

Dick: Waterboarding is not torture; torture requires inflicting “pain that is difficult for the individual to endure and is of an intensity akin to the pain accompanying serious physical injury,” and waterboarding is not painful in that way.⁹

(12) Anne: Lia Thomas is a woman.

Brie: Lia Thomas is not a woman; to be a woman, you have to have two X chromosomes.¹⁰

⁷ Plunkett and Sundell, “Disagreement and the Semantics of Normative and Evaluative Terms,” 17.

⁸ Plunkett and Sundell, “Disagreement and the Semantics of Normative and Evaluative Terms,” 24.

⁹ The quoted passage comes from Jay S. Bybee, “Memorandum for John Rizzo, Acting General Counsel of the Central Intelligence Agency: Interrogation of Al Qaeda Operative,” August 1, 2002, 9–10, <https://www.justice.gov/sites/default/files/olc/legacy/2010/08/05/memo-bybee2002.pdf>.

¹⁰ Lia Thomas, a swimmer, was the first openly transgender athlete to win an NCAA Division I national championship.

These disputes are certainly not trivial; matters of great consequence hinge on how they are resolved. But this does not show that the locus of disagreement is not metalinguistic. How we use words like “woman” and “torture” can matter a great deal because of the role these words play in laws and social practices.¹¹

Plunkett and Sundell suggest that disagreements like (11) and (12) amount to disagreements about what concept best fills a certain functional role.¹² However, this way of dispelling the triviality objection may be problematic for them, given their other commitments. Once it is conceded that the disagreeing parties have a common conception of the functional role “torture” plays in our society, one could argue that both parties mean by “torture” *whatever plays this functional role in our society*. If nontrivial disagreements can always be conceived as disagreements about what best fills a functional role, then one way to save the argument from disagreement would be to let this functional role fix the shared meaning.

I won’t pursue this issue further here, because I do not think that responding to the triviality worry requires this functionalist gloss on metalinguistic disagreements. Instead, I want to focus on Assertoric Compatibility: the idea that the two parties are attaching different meanings to their words, and therefore making compatible assertions. I think there are two strong reasons for doubting that the speakers make compatible assertions in these cases: the problem of missing agreement (discussed in section 5.3) and the problem of inscrutable content (discussed in section 5.4).

5.3 Missing Agreement

When speakers are making compatible assertions and disagreeing only in the normative attitudes they express, it should be possible for them to

¹¹ Even the decision whether to apply “athlete” to horses has ethical ramifications. If race horses are athletes, we can conceptualize their suffering as self-sacrifice for their own goals; critics charge that this allows us to rationalize what amounts to mistreatment. A. N. McLean and P. D. McGreevy, “Ethical Equitation: Capping the Price Horses Pay for Human Glory.” *Journal of Veterinary Behavior: Clinical Applications and Research* 5, no. 4 (2010): 203–9. I am thankful to Sofia Jeppsson for the reference.

¹² Plunkett and Sundell, “Metalinguistic Negotiation and Matters of Language.”

express *agreement* as well as disagreement. And that is exactly what we find in some of our earlier motivating examples:

(13) Anne: Your wife is very beautiful.

Bjorn (a): No, we're not married.

Bjorn (b): Yes, Sarah is beautiful. But we're not married.

(14) Ani: She is good-looking.

Beth (a): No, she is *gorgeous*.

Beth (b): Yes. But you're understating; in fact, she is *gorgeous*.

(15) Abe: The Sioux used to live on this land.

Beth (a): No, the *Lakota* used to live here.

Beth (b): Yes, but we should call them *Lakota*.

In these cases, either continuation would be felicitous, and that is consistent with the claim that the disagreement does not concern the main asserted content.

We see the same pattern in (2):

(16) American: The burgers come with chips.

Brit (a): No, they come with *crisps*.

Brit (b): Yes. But around here, you should call them *crisps*.

But expressions of agreement do not seem felicitous in (3) or (4):

(17) Host: Secretariat was one of the best athletes of the last century.

Caller: ??Yes, I agree. But you shouldn't call a racehorse an *athlete*; we ought to reserve that word for humans.

(18) Kamala: People who make \$400,000 a year are rich.

Donald: ??Of course. But I think "rich" should be used only for people with yachts and private planes, and you need a lot more for that.

I think this is a good indication that the speakers do not take themselves to be asserting compatible contents in these cases.

5.4 Inscrutable Contents

In cases like (16), where agreement is felicitous, the speakers are clear about what they are agreeing to and can make this explicit:

(19) American: The burgers come with chips.

Brit: Yes, they do come with thin fried potatoes. But around here you should call them *crisps*.

Things are very different in the other cases. In (17), there is no particular meaning that Caller is in a position attribute to Host's word "athlete." As a result, even if he wanted to agree, there would be no natural way to say what he is agreeing to:

(20) Host: Secretariat was one of the best athletes of the last century.

Caller: ??Yes, I agree that _____. But you shouldn't call a race-horse an *athlete*; we ought to reserve that word for humans.

Similarly, in (18), there is no natural way for Donald to rephrase the content he would be agreeing to, in agreeing with Kamala:

(21) Kamala: People who make \$400,000 a year are rich.

Donald: ??Yes, you're right, they are _____. But I think "rich" should be used only for people with yachts and private planes, and you need a lot more for that.

Plausibly, agreement is blocked in these cases because the content asserted is inscrutable.

We can develop this thought into an independent argument. Asserting that *p* requires intending that it be common ground that one is asserting that *p*. But the speakers in our examples cannot reasonably intend this, if Plunkett and Sundell's analysis is correct.

Let me first say a bit to defend my premise. I assume that asserting is a case of speaker-meaning, in the sense analyzed by Paul Grice.¹³ It is fundamentally a communicative act, and as such, it essentially involves an intention to get one's audience to recognize what speech act one is performing. The reason I can't, right now, assert that it is past 6 p.m. by saying "The ants have crawled past the bottom of the sink" is not that

¹³ H. P. Grice, "Meaning," *Philosophical Review* 66, no. 3 (1957): 377–88.

these words are fundamentally unsuitable for making this assertion—after all, we might adopt the convention of thinking of the big hand of the clock as a group of ants, and the clock face as a sink—but that, without the right kind of setup, I can't expect you to associate this message with those words, and therefore I can't intend that you grasp what I am trying to assert. (I assume here that it is impossible to intend things that you take to be essentially impossible.) For similar reasons, I can't assert that a particular book is red by saying "That is red" unless I've given you sufficient clues to figure out what I mean to refer to by "That." Otherwise, I cannot expect, and cannot rationally intend, for it to be common ground what I am asserting.

Another way to motivate the premise starts from Robert Stalnaker's influential account of assertion. On Stalnaker's view, assertion takes place against the background of an evolving "common ground" of assumptions accepted for the purposes of the conversation. An assertion can be thought of as (inter alia) a proposal to add its content to the common ground—a proposal that can be accepted or rejected. As Stalnaker notes, 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. For this reason, Stalnaker takes it to be a fundamental pragmatic principle that it be common ground what proposition is being asserted.¹⁴

One of the principal ways we can achieve this kind of common ground about the contents of our assertions is by piggybacking on a public meaning. I assert, "Someone discovered a new quasar today," and even though I can't say much about what "quasar" means, I can expect you to recognize that I'm announcing the discovery of one of those things to which "quasar," in its standard public meaning, applies. But according to Plunkett and Sundell's account of cases like (3) and (4), one or both parties are using a word to mean something different from its public meaning—what they think the public meaning *ought* to be, not what it is. How, in that case, can they expect the other party will cotton on to what is being asserted?

¹⁴ Robert Stalnaker, *Context and Content: Essays on Intentionality in Speech and Thought* (Oxford University Press, 1999), 88, discussed earlier in chapter 1, section 1.6.

Let's look at the "athlete" case, (3). Given Assertoric Compatibility, we can think of the two parties as asserting compatible contents:

(22) Host: Secretariat was one of the best athletes_H of the last century.

Caller: No, Secretariat is not one of the best athletes_C of the last century.

where "athletes_H" and "athletes_C" are disambiguated because they express different the concepts of *athlete* Host and Caller are using. But how is Caller supposed to grasp the content of "athletes_H," and how is Host supposed to grasp the content of "athletes_C"? Caller knows that "athletes_H," unlike "athletes_C," can apply to horses. But knowing that is not enough to grasp its content. Can it apply to dogs? Lizards? Snails? Robots? Caller knows that it is different from "athletes_C" in this one respect, but that is not enough to understand what Host is asserting. Similarly, Host knows that "athletes_C" excludes nonhumans, but that is not enough to grasp its content. There are many possible concepts that include paradigm athletes and exclude nonhumans. If the two parties could both defer to a public meaning, then we wouldn't need to worry about this indeterminacy, but then there would be a shared content that one was asserting and the other denying, contradicting Assertoric Compatibility.¹⁵

Similarly, if we disambiguate senses of "rich" in (4), we get something like

(23) Kamala: People who make \$400,000 are rich_K.

Donald: No, people who make \$400,000 are not rich_D.

¹⁵ The argument is very close to one in Tyler Burge, "Individualism and the Mental," *Midwest Studies in Philosophy* 4, no. 1 (1979): 94. Arguing against a reinterpretation of his "arthritis" thought experiment that combines the idea that the speakers are asserting idiosyncratic contents with the idea that the disagreement between them is metalinguistic, Burge objects: "Suppose we are to reinterpret the attribution of his erroneous belief that he has arthritis in the thigh. We make up a term 'tharthritis' that covers arthritis and whatever it is he has in his thigh. The appropriate restrictions on the application of this term and of the patient's supposed notion are unclear. Is just any problem in the thigh that the patient wants to call 'arthritis' to count as tharthritis? Are other ailments covered? What would decide? The problem is that there are no recognized standards governing the application of the new term."

We can imagine that the difference between “rich_K” and “rich_D” is that they put the threshold for being rich in different places. It is common knowledge that the threshold for “rich_K” is under \$400,000 and the threshold for “rich_D” is considerably over \$400,000. But knowing this is not enough to grasp “rich_K” and “rich_D” since there are many places to put the thresholds that meet this constraint. Donald can’t know exactly what threshold governs Kamala’s use of “rich,” and Kamala can’t know exactly what threshold governs Donald’s. But then it seems that Donald and Kamala can’t intend for it to be common ground what they are asserting.

The problem doesn’t arise for (2), because here, it is plausible that American can intend Brit to recognize that he is using “chips” in the American way (for thin fried potatoes) and Brit can intend American to recognize that he is using “chips” in the British way (for thick fried potatoes). Here, it’s common ground that there are just two reasonable options, and it’s also common ground what those options are. That’s probably not the case in (3), where the parties aren’t just selecting among a fixed set of standard meanings, but trying to innovate. And it’s *definitely* not the case in (4). The “chips” case, then, is rather special, and not much like the philosophically interesting cases of metalinguistic negotiation. Plunkett and Sundell seem to miss this difference. Indeed, they rely on the claim that “the ‘athlete’ case and the ‘chips’ case have precisely the same structure” to fend off an externalist challenge to Assertoric Compatibility.¹⁶

We should concede, I think, that Plunkett and Sundell’s analysis works well for (2). In this case, it is plausible that we have both Assertoric Compatibility and Expressive Incompatibility. But in cases like (3) and (4), and certainly in philosophically interesting cases involving words like “woman” or “torture,” Assertoric Compatibility is implausible—unless we suppose that the participants in these conversations are under a serious misconception about what they are doing.

5.5 Confused Speakers

As it turns out, Plunkett and Sundell are willing to bite this bullet: They are prepared to attribute massive error to speakers about what they are

¹⁶ Plunkett and Sundell, “Disagreement and the Semantics of Normative and Evaluative Terms,” 27.

asserting.¹⁷ The parties to a metalinguistic negotiation, they say, may not realize that they are asserting compatible contents. A dispute like (3) or (4) might feel to them much like an ordinary, first-order disagreement.

Such an error theory demands an explanation: How do the speakers fail to see that they are asserting compatible contents? Plunkett and Sundell say this:

When speakers involved in a dispute *take themselves to be at odds* there is little reason to think that they have fine-grained intuitions about the modes of communication with which they express their disagreements, or about subtle theoretical distinctions amongst equally substantive kinds of disagreements. Ordinary speakers do not have intuitions of the kind that could make subtle distinctions about the semantics/pragmatics distinction. How could they, when theorists are in many cases divided about which aspects of meaning belong on which side of the divide?¹⁸

We can certainly agree that ordinary speakers do not have sophisticated beliefs about the contributions of semantics and pragmatics in determining the asserted content. But the questions Plunkett and Sundell are taking ordinary speakers to be confused about are simply: What is being asserted? What is being denied? And are they the same thing? These questions will arise whether the asserted contents correspond to the literal meanings of the sentence used or result from some kind of pragmatic enrichment. Now, we *could* take such questions to be beyond the ken of ordinary speakers. But if we do that, we have to reject the view that asserting something requires intending that it be commonly known what one is asserting. If ordinary speakers have no grip on what has been asserted and what is being denied, then we can't think of them as engaged in the sort of rational practice Stalnaker describes, which requires them to keep track of what has been asserted and what is part of the common ground.

I am skeptical, then, that attributing this sort of error to ordinary speakers is plausible. But let's concede for the sake of argument that speakers might have an inaccurate view of these matters—one that

¹⁷ Plunkett and Sundell, "Disagreement and the Semantics of Normative and Evaluative Terms," 23ff; Plunkett and Sundell, "Metalinguistic Negotiation and Speaker Error."

¹⁸ Plunkett and Sundell, "Disagreement and the Semantics of Normative and Evaluative Terms," 23, emphasis in the original.

we theorists could clear up. Then a more troubling objection looms, for Plunkett and Sundell's account *depends* on the speakers' recognizing that they are using their words with different meanings, and therefore asserting compatible contents.

Consider why it was plausible to suppose that in using "athlete" as they do, Host and Caller express incompatible normative proposals about how the word should be used. If Host and Caller are *openly and intentionally* using "athlete" with different meanings, then we can see why their speech acts express incompatible normative metalinguistic proposals. Each is openly presenting himself as doing something that could be right only if a certain metalinguistic proposal were accepted; each intends others to recognize this and infer that he endorses this proposal.¹⁹ So here we have a nice explanation of why a normative attitude is expressed, but it depends on the speakers' being aware that they are using the words with different meanings. If, instead, they take themselves to mean the same thing by "athlete" and present themselves that way to each other, why should their speech acts express incompatible proposals about how to use "athlete"?²⁰

In positing speaker ignorance here, Plunkett and Sundell are undermining their own story. They can't have it both ways. If they want to explain Expressive Incompatibility, then they need to suppose that the speakers in cases of metalinguistic negotiation recognize that they are asserting compatible contents (Assertoric Compatibility). But then they face the objections from inscrutable content and missing disagreement. If they seek to avoid these objections by positing speaker error, then they can no longer explain Expressive Incompatibility.

5.6 Mixed Cases

Plunkett and Sundell's account splits the impact of a speech act into two parts: a descriptive part, which is asserted, and a normative, meta-

¹⁹ "For *S* to express an attitude is for *S* to *R*-intend the hearer to take *S*'s utterance as reason to think *S* has that attitude." Kent Bach and Robert M. Harnish, *Linguistic Communication and Speech Acts* (MIT Press, 1979), 15.

²⁰ In their initial presentation of their view, before embracing speaker ignorance, Plunkett and Sundell say that it is *via* asserting compatible propositions that the speakers advocate for different meanings ("Disagreement and the Semantics of Normative and Evaluative Terms," 15). That seems just right.

linguistic part, which is expressed in making the assertion. But there are cases in which this splitting is not possible, and they pose a further challenge to the account.

Suppose that Donald and Kamala both know a man named Fred, but neither is sure what the other knows about Fred. It is common ground between them that Fred makes less than \$2 million per year, but aside from this, Kamala doesn't know what Donald knows about Fred's income. (Nor does Donald think she does.) It is also common ground that Donald thinks that "rich" can permissibly be applied to people who make \$2 million per year, but aside from this, Kamala doesn't know how rich Donald thinks you have to be to count as "rich." (Nor does Donald think she does.)

Suppose that in this context, Donald asserts:

(24) Fred is not rich.

Has Donald thereby expressed any normative view about where the threshold for "rich" should lie? It's hard to see why he would have. Even someone with Kamala's views about who counts as rich might assert (24), if they knew independently that Fred made \$50,000 per year. If it were common ground that Fred makes \$400,000 a year, then Donald would be expressing a normative view—the view that "rich" should not be applied to people who make \$400,000 a year. But by hypothesis, the only thing that is common ground about Fred's salary is that it is less than \$2 million a year.

It seems, then, that Donald hasn't expressed any normative meta-linguistic attitude in asserting (24). But he also has not asserted any descriptive proposition, for to do this, he would have to have ruled out some possible states of affairs, and he hasn't. His assertion doesn't rule out any possible salary for Fred that is not already excluded by the common ground. Even if Donald thought Fred made \$1,999,999, he might (for all that is assumed in common) think this isn't enough to count as "rich."

So, in cases like (24), we cannot factor the impact of the speech act into two separate components—a purely normative metalinguistic update and a purely descriptive nonmetalinguistic update.

The problem here is analogous to a problem facing the cruder forms of expressivism in meta-ethics. The expressivist wants to say that

(25) Jean Valjean did wrong in stealing the silver

conveys the information that Jean Valjean stole the silver and expresses a negative attitude towards that act. But what about, say,

(26) Jean Valjean did something wrong while staying with Monseigneur Bienvenu.

If it's not common ground what Jean Valjean did, then (26) doesn't express any particular attitude towards any particular act. (It could have been uttered by someone who thinks that Jean Valjean offered his host money, for example, and disapproves of that, but approves of stealing silver.) And if it's not common ground what the speaker's attitudes are, then asserting (26) conveys no information about what Jean Valjean did (or at any rate, no information we could convey without using the word "wrong").

To handle this kind of mixed case, Allan Gibbard proposed combining the factualist and nonfactualist elements of the assertion into a single content.²¹ However, if proponents of metalinguistic negotiation do this, they are giving up Assertoric Compatibility, and their analysis is no longer a threat to the argument from disagreement.

5.7 A Diagnosis

I think that these objections to Plunkett and Sundell's account are compelling. But one thing I learned from John McDowell is that you can't kill a philosophical theory just by giving arguments against it. If its proponents think that the only alternatives to their theory are even more implausible, then they will defend it no matter how many bullets they have to bite. What one needs is to offer them a plausible alternative—and this sometimes requires getting them to give up a picture that structures their space of alternative views.

I think that Plunkett and Sundell's account has seemed attractive to philosophers mainly because they think the alternatives are even worse. So, let's think about what the alternatives are, if we want to say that in

²¹ Allan Gibbard, *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Harvard University Press, 1992); Allan Gibbard, *Thinking How to Live* (Harvard University Press, 2003).

disputes like (3) or (4), there is a single content that is being asserted by one party and denied by the other.

One option is to complete what Plunkett and Sundell call the “shared meaning task”: to come up with a definition, in noncircular terms, that all the disputing parties could accept—“a meaning that, despite the systematic variation in usage, allows for the conceptual coherence of their respective views of what falls into the extension of the term.”²² As they note,

This is a difficult task: the theorist must thread a very fine needle, constructing an analysis which is substantive enough to differentiate the expression under analysis from other, nearby expressions, while simultaneously being so devoid of specific content as to plausibly be the meaning shared by speakers who apply the term in systematically differing ways.²³

Although meta-ethicists have sometimes claimed to have pulled off this feat for contested ethical terms, it is difficult to see how it could be done for “rich” in (4) or “athlete” in (3). And its requirement of an analysis or definition should worry us: even if we think our grasp of some terms can be made explicit in a noncircular definition, this can’t be the case for *every* term. Ludwig Wittgenstein’s remarks on “game,” as well as the sad history of failures to give informative analyses of philosophical terms, might make us worry that the shared meaning task is hardly ever solvable.²⁴

The remaining alternative, it may seem, is something we might call “blind externalism.” This is the view that the community’s usage of a term, in conjunction with relevant facts about the context in which it being used, fixes its content in a way that is entirely opaque to us. This is more or less the view Timothy Williamson endorses about the contents of vague words: Somehow, usage fixes a meaning for “young” that determines an exact second that was the last second of my youth,

²² Plunkett and Sundell, “Disagreement and the Semantics of Normative and Evaluative Terms,” 25.

²³ Plunkett and Sundell, “Disagreement and the Semantics of Normative and Evaluative Terms,” 25.

²⁴ See Ludwig Wittgenstein, *Philosophical Investigations*, 2nd ed. (Blackwell, 1958), sections 66–71; Timothy Williamson, *The Philosophy of Philosophy* (Blackwell, 2007), chapter 4.

and how it does so we cannot say. "Meaning may supervene on use in an unsurveyably chaotic way."²⁵

I call this "blind" externalism to distinguish it from the sorts of externalism Hilary Putnam and Tyler Burge convinced nearly everyone to accept.²⁶ "Water" has a meaning that applies only to H₂O, not to indistinguishable liquids with very different chemical compositions, and it had this meaning even when people thought of water as an element, because they intended to use "water" as a term for a kind of stuff—the kind found in our lakes, rivers, and seas. Here, the way in which the reference of "water" is determined is not opaque to us. Similarly, Burge's character who thinks that one can have arthritis in the thigh can still mean *inflammation of the joints* by arthritis, because his dominant intention is to mean what everyone else means, and to defer to experts (doctors) about the proper use of "arthritis." Again, it is clear how the meaning is determined. But these standard externalist stories are not available for "athlete" or "rich." These are not substance words, and there are no experts here whose opinions we are disposed to defer to as settling their content. So the only kind of externalism that seems plausible is blind externalism.

Blind externalism is hard to accept. If nobody knows how the contents of "rich" and "athlete" are fixed, then shouldn't we be much more hesitant in our use of these words? Shouldn't Donald and Kamala agree that one of them is right and one of them is wrong, but that nobody knows how to determine which is which? But in that case, how can they be so confident in making their assertions?²⁷ It seems that if blind externalism is correct and there are facts about the contents of these words that are mysteriously determined by use, then these facts, which are utterly inaccessible to speakers, can play no real role in explaining language use.

If the options are blind externalism or the metalinguistic negotiation account, it would not be unreasonable to prefer the latter. Intuitively,

²⁵ Timothy Williamson, *Vagueness* (Routledge, 1994), 209.

²⁶ Hilary Putnam, "The Meaning of 'Meaning'," *Minnesota Studies in the Philosophy of Science* 7 (1975): 131–93; Burge, "Individualism and the Mental."

²⁷ Williamson himself appeals to our semantic ignorance about the boundaries drawn by vague terms to explain our hesitation to apply vague terms in borderline cases. That explanation would be threatened if we embraced blind externalism in cases where we find no similar hesitation.

there is a certain flexibility to words like “rich” and “athlete.” If speakers want to talk of “athlete” as applying to horses, then that’s OK: It need not impede communication or any of the other ends of language, so long as others go along with it. If they want to use “rich” in such a way that even people who make \$400,000 a year are not rich, then that’s OK too—perhaps it is useful for some purposes. Neither person need be making a mistake. The metalinguistic negotiation story makes sense of this, while blind externalism makes it mysterious. The thought that Host and Caller are not contradicting each other may be odd, but the thought that some unknown factor somehow settles whether “athlete” applies to horses seems even odder.

5.8 An Alternative

Fortunately, there is a third option. Like Plunkett and Sundell’s account, it recognizes that speech acts may involve competing normative proposals about how words should be used. But unlike their account, it takes disputes like (3) and (4) to concern a common asserted content.

To say both these things, we need to revise standard assumptions about the contents of assertions, along the lines proposed in chapter 3. We need to drop the view that having a common content involves agreeing on an intension, a definition, or a particular means of reference fixation. As the conversation proceeds, the participants can sharpen up their plans about what is to count as *rich*—for example, deciding that only those who can afford to fly first class count as rich—but this sharpening should not be conceived as changing what they *mean* by “rich.” Instead, we should embrace Hilary Putnam’s slogan that “meanings have an identity through time but no essence.”²⁸

This may seem a bit mysterious, so let me sketch a formal model that would make sense of this picture, based on the pragmatics of chapter 3 and the panvariabilism introduced in chapter 4. We’ll treat words like “athlete” and “rich” as *variables*, which are given their intensions by an assignment function v .²⁹ And we’ll think of the contents of declarative sentences as sets of world/assignment pairs. For example, the content of

²⁸ Hilary Putnam, *Representation and Reality* (MIT Press, 1988), 11.

²⁹ In standard semantic theories, assignment functions assign *objects* to first-level variables. Here we are supposing that variables can come in various *types*; for example, “athlete” has type $(w, (e, t))$ and so the assignment function will assign it a value of

(27) People who make \$400,000 a year are rich,

is the set of pairs (w, v) such that the intension that v assigns to “rich” applies in w to people who make at least \$400,000 a year.

You might ask: Why separate these two components, w and v ? Why not consider the assignment of an intension to “rich” a part of w , so we can dispense with v ?³⁰ Answer: Because there isn’t a fact about what the intension of “rich” is. Even if all the facts of the world are settled, it isn’t settled what intension “rich” has at that world. This isn’t meant as a deep metaphysical claim, but as an articulation of the view that ordinary speakers have (as discussed previously in chapter 3, section 3.1). When they think that something is a factual matter, though unknown, they will use expressions of epistemic uncertainty, for example:

(28) That pole might be more than 10 meters tall.

(29) That pole is probably more than 10 meters tall.

But in contested cases involving “athlete” or “rich,” this way of talking seems inapt:

(30) ??People who make \$400,000 might be rich, but maybe Donald is right that you need a private plane.

(31) ??Secretariat is probably not an athlete, but Caller might be right that he is.

Instead, it seems appropriate to express our hesitation using the language of decision:

(32) We could count people who make \$400,000 as rich, but we could also say you need a private plane to be rich.

(33) We could count Secretariat as an athlete, but we ought not to.

On the other hand, this language would not be apt when talking of the height of the poles (except in very special cases, where there’s indeterminacy about how to measure the height):

this type: a function that takes a world and gives you a function from entities to truth values. If we use a standard degree semantics for “rich,” then the assignment function won’t actually assign a value directly to the positive form of “rich,” which is actually a compound expression consisting of a null morpheme *POS* plus the root meaning of “rich” (see chapter 4, section 4.2) and possibly an implicit domain restriction like “for an American.” I’ll ignore this complexity to simplify presentation of the view.

³⁰ Compare Stalnakerian diagonalization, discussed in chapter 1, section 1.6.

- (34) ??We could count that pole as being more than 10 meters high.
 (35) ??We ought to count that pole as being more than 10 meters high.

All of this suggests that, at least as ordinary speakers view things, there is a fact about the height of the pole, but what counts as *rich* is (at least in part) a matter of decision. The assignment component of our pairs (v) represents the nonfactual, decisional aspect (how we plan to use words) while the world component (w) represents the factual aspect.

A content—a set of such pairs—can be viewed as a joint constraint on our plans for using words and our views about what the world is like. Accepting a content in which w_1 never appears in the world component of a pair is *ruling out* a factual state of affairs, w_1 . Similarly, accepting a content in which v_1 never appears in the assignment component is *ruling out* an assignment of meaning, v_1 . But it is also possible for a content to constrain *combinations* of meaning-assignments and worldviews without ruling out any particular meaning-assignments or worldviews. For example, accepting the content

$$\{(w_1, v_2), (w_2, v_1)\}$$

doesn't rule out either world (w_1 or w_2) or either assignment (v_1 or v_2), but it *does* constrain us not to rule out *both* w_1 and v_1 .

Recall that Plunkett and Sundell's account faced a problem with "mixed cases," such as disjunctions in which one disjunct contains the contested word but the other does not (section 5.6). We could not explain in such cases how a normative metalinguistic attitude would be expressed, since it is possible to accept the disjunction without having any particular normative attitude. The root problem is that Plunkett and Sundell's account attempts to separate the normative, metalinguistic aspect of assertion from the descriptive, worldly aspect, positing a separate attitude for each. In our model, by contrast, we have a single content whose acceptance involves both normative and doxastic commitments.

For example, accepting

- (24) Fred certainly isn't one of the rich and happy people.

rules out combining a doxastic view that takes Fred to be happy and to have a salary of \$X a year with a decision to put the threshold for counting as "rich" below \$X a year. Whether this has the force of a normative,

metalinguistic proposal or a factual claim depends on what is common ground when the assertion is made. If the common ground accepts that Fred makes \$400,000 a year but rules out thresholds for “rich” above \$300,000, then (24) may simply convey the factual information that Fred is not happy. On the other hand, if the common ground accepts that Fred is happy and makes \$400,000 a year, then (24) may express a plan to use “rich” with a higher threshold than \$400,000. And if the common ground makes neither assumption, then asserting (24) may have mixed doxastic and planning import. The important thing to see is that its import depends on the common ground and is not built into its content.

Indeed, the content of “rich” on this view is exceedingly thin: To simplify somewhat, it is a function that maps an object o to the set of world/assignment pairs (w, v) such that o falls into the extension of “rich” at w .³¹ This content is entirely neutral as to where the threshold for being rich lies. This neutrality is what allows us to have a common content across people (and conversations) with different views about how wealthy one must be to count as “rich.”

On this view, there seems to be only a *formal* difference between the meaning of “rich” and the meaning of, say, “tall.” Isn’t that a *reductio* of the view? But it is important to recognize that we do not come to conversations armed *only* with our knowledge of the linguistic meaning of these words. Communication always takes place against a background of common assumptions, a common ground. In addition to factual assumptions (for example, that the United States has a president), the common ground will include shared plans about the use of words like “rich” and “tall.” That is what allows us to start using these words in a conversation without an elaborate setup. As the conversation continues, these plans may become more determinate. But we can hold on to a topic—a proposition whose truth is at stake—while these shared plans are refined. What makes it the case that there is a single content at stake in a conversation, across such changes, are not facts about the classificatory dispositions of the individual participants (or even the larger society), but the participants’ intentions to be agreeing and disagreeing with each other.

³¹ This is a simplification, because I’ve left out some details that would be needed in a proper semantics for gradable adjectives—for example, the relativity to a contextually supplied reference class (see section 4.2 in chapter 4 and n29 in this chapter).

Consider an analogy to variables in mathematics. When faced with a system of equations like

$$y = x^2$$

$$y = z + 3,$$

it is senseless to ask in virtue of what the y in the first equation has the same value as the y in the second. This is *stipulated* to be so; it is part of the convention of repeating a variable name that all like-named variables in the same group of equations get the same value. The identity of denotation of the first y and the second y is not grounded in independent facts about what these occurrences denote; rather, it is presupposed from the beginning that these occurrences mean the same thing, and this constrains the facts about what they denote. I want to say that the same thing happens with “rich” in (4) and “athlete” in (3).

This is not to say that we have the convention that all orthographically identical words get the same meaning. We have homonyms in natural languages, and it is even possible to have a conversation in which different occurrences of “rich” are not considered the same variable. Sameness here is determined not by orthography, but by the participants’ intentions to agree and disagree.

So, that’s a sketch of an alternative picture.³² I believe this sort of view captures everything that is plausible about the metalinguistic negotiation account, without being committed to any of its implausible consequences. It is tempting, indeed, to think of this view as a *version* of the metalinguistic negotiation account, and there are some hints that Plunkett and Sundell would think of it this way. At one point, they seem to class Chris Barker’s view as a version of the metalinguistic negotiation account, even though it posits a single content that is asserted by

³² Views that are similar in some respects have been proposed by Chris Barker, “Negotiating Taste,” *Inquiry* 56, nos. 2–3 (2013): 240–57, Alexander W. Kocurek, Ethan Jerzak, and Rachel Etta Rudolph, “Against Conventional Wisdom,” *Philosophers’ Imprint* 20, no. 22 (2020): 1–27; Alexander Kocurek, “Verbal Disagreement and Semantic Plans,” *Erkenntnis* 90 (2023): 1061–94; Andrés Soria-Ruiz, “Is Metalinguistic Usage a Conversational Implicature?” *Topoi* 42 (2023): 1027–38. Kocurek et al., “Against Conventional Wisdom,” 15, note the connection between their view and the expressivist account of vagueness given in John MacFarlane, “Lecture III: Indeterminacy as Indecision,” *Journal of Philosophy* 117, no. 11/12 (2020): 643–67, which appears with revisions as chapter 3 of this book.

one party and denied by the other.³³ However, if this is what metalinguistic negotiation amounts to, it poses no threat to the argument from disagreement. Indeed, because it grounds facts about content identity on intentions to agree and disagree, it presupposes the validity of the argument from disagreement.

5.9 Half-Measures

It is useful to compare the account I have offered with some other accounts of metalinguistic negotiation, which do better than Plunkett and Sundell's account but retain too much of its structure to be fully satisfying solutions.

Single-Scoreboard Semantics

The most problematic assumption of Plunkett and Sundell's model is that the two parties are asserting compatible contents. That assumption gives rise to the problems of missing agreement (discussed in section 5.3), inscrutable contents (discussed in section 5.4), and confused speakers (discussed in section 5.5). Responding to similar objections against contextualist accounts of "knows," Keith DeRose proposed a way for a contextualist to avoid the implication that contending speakers assert compatible contents.³⁴ The central idea is that "knows" is contextually sensitive to an "epistemic standards" parameter on a shared "scoreboard." (The metaphor and the general strategy comes from David Lewis.)³⁵ When the dogmatist says, "Moore knows he has hands" and the skeptic says, "Moore doesn't know he has hands," they are contending about the truth of a single proposition, the proposition that

- (36) Moore's epistemic position with respect to the proposition that he has hands exceeds the standard set by the shared scoreboard.

³³ Plunkett and Sundell, "Metalinguistic Negotiation and Speaker Error," 153; the view they are discussing is that of Barker, "Negotiating Taste."

³⁴ Keith DeRose, "Single-Scoreboard Semantics," *Philosophical Studies* 119 (2004): 1–21.

³⁵ "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8, no. 3 (1979): 339–59.

The dogmatist asserts this proposition and the skeptic asserts its negation. The word “knows” is context-sensitive because its intension varies with a feature of the context—the epistemic standards parameter of the shared scoreboard. But because this feature of the context is the same for both speakers, the intension is also the same for both of them. Both speakers may be aware that it is contentious what the shared standard is, but by acting in a way that makes sense only if the scoreboard is a certain way, they pressure their interlocutors to make it so, through what Lewis calls “accommodation.”³⁶

Applied to a case like (4), the single scoreboard model would say that, instead of asserting the compatible propositions

(37) People who make \$400,000 are rich_K

(38) People who make \$400,000 are not rich_D,

Donald and Kamala express the incompatible propositions

(39) People who make \$400,000 are rich_S

(40) People who make \$400,000 are not rich_S,

where “rich_S” means “rich by the standards established by our shared conversational scoreboard.”

While this approach allows us to avoid several of the objections we raised against Plunkett and Sundell’s view, it does not help with the problem of mixed cases (discussed in section 5.6). It presupposes that the common ground can be factored into two parts—a factual common ground (representable as a set of worlds or a set of propositions) and a set of nonfactual parameters that affect context-sensitive expressions. (In chapter 2, section 2.3, we called this a “bifurcated common ground.”) Cases like the dispute between Donald and Kamala constitute disputes about what the nonfactual component should be, and cases like

(41) Andy: The median US salary is less than \$50,000.

Beatrice: The median US salary is at least \$50,000.

constitute disputes about what the factual component should be. But there are plenty of cases that fall between these two stools. For example, if it isn’t common knowledge between two parties which thresholds the other accepts for counting as rich, and it isn’t common knowledge how much Fred earns, then a conversation like

³⁶ See chapter 2, section 2.3.

(42) Ava: Fred is rich.

Ben: No, Fred isn't rich.

cannot be understood as a dispute about the threshold or as a dispute about the world (Fred's salary)—or as both disputes put together. Given what Ava and Ben know in common, all of the following situations are possible:

- Ava and Ben have the same beliefs about Fred's salary but differ on what the threshold for "rich" should be.
- Ava and Ben have the same views about what the threshold for "rich" should be but differ about Fred's salary.
- Ava and Ben differ both about what the threshold for "rich" should be and about Fred's salary.
- Ava or Ben (or both) don't have fixed views about what the threshold for "rich" should be *or* about Fred's salary, but they consider these questions linked in a way that ensures the truth of their claim.

Discourse Contextualism

Alex Silk's "discourse contextualism" closely resembles the single-scoreboard view:

The lesson is this: *by acting in such a way that is appropriate only if the context is a certain way, one can implicitly propose that the context be that way.* If the other party accommodates by proceeding in like manner, it can become taken for granted that the context is that way. If she doesn't, this can lead to negotiation over the state of the context. Crucially this can all happen without explicitly raising the issue of what the context is like. I suggest that the linguistic case—the case of linguistic action, discourse, and interpretation—is a special instance of these phenomena.³⁷

Like Plunkett and Sundell, and like DeRose, Silk emphasizes that his account requires no changes to standard contextualist semantic and pragmatic assumptions.³⁸ In particular, we do not need to innovate in our account of asserted contents, which can still be represented as sets of possible worlds.

³⁷ Alex Silk, *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics* (Oxford University Press, 2016), 25, emphasis in the original.

³⁸ Silk, *Discourse Contextualism*, 3.

Unlike Plunkett and Sundell, and unlike DeRose, Silk recognizes the existence of mixed cases, which he calls “hybrid uses”:

In addition to purely descriptive uses and purely normative uses of deontic modals, there are what we can call *hybrid* uses. Hybrid uses are uses that are neither purely descriptive nor purely normative. They rule out certain combinations of norms and nonnormative facts. Hybrid uses constitute the typical case.

He gives the following example:

Consider an utterance of

(43) Sally must give 10 percent of her income to the poor.³⁹

Suppose that prior to the speaker’s utterance it is compatible with the common ground that Sally’s giving 10 percent of her income to the poor would produce 125 or 175 utils of happiness. And suppose that some satisficing version of utilitarianism has been accepted, but it hasn’t been settled precisely what the relevant threshold is. Normative thresholds that require Sally to produce 100, 150, or 200 utils are all compatible with what has been accepted thus far.⁴⁰

This is exactly the kind of case that we saw to be problematic for views that bifurcate the common ground into a worldly and nonfactual parts, because the update proposed by (43) cannot be understood either as an update to the factual part or as an update to the nonfactual part. Silk sees the problem clearly: “Updating with [(43)] in this context doesn’t settle either on the relevant body of norms or on the nonnormative facts about the consequences of the relevant acts, but it does rule out certain combinations thereof.”⁴¹

To deal with the problem, Silk employs Stalnakerian diagonalization (previously discussed in chapter 1, section 1.6).⁴² He assumes that each possible world in the context set determines an “abstract context that represents the conversational situation in *w*,” and this abstract context determines a value for the normative threshold. The update proposed by

³⁹ I have renumbered Silk’s example and repeated it to avoid confusion with our own numbered examples. The original just says “(5).”

⁴⁰ Silk, *Discourse Contextualism*, 134, emphasis in the original.

⁴¹ Silk, *Discourse Contextualism*, 135, with the sentence number changed.

⁴² He doesn’t call it “diagonalization” here, but earlier he remarks on the similarity between his story about updates and Stalnaker’s (Silk, *Discourse Contextualism*, 37n11).

(43) can then be represented as a set of possible worlds: the set of worlds w such that the number of utils Sally produces in w by giving 10 percent of her income to the poor meets the normative threshold specified by the abstract context determined by w .⁴³

It is a strength of Silk's account that it recognizes and tries to meet the challenge posed by mixed cases. But Silk's answer to that challenge requires us to assume that a possible world determines an abstract context, and thereby a value for the normative threshold. To assume this is to make our ambivalence about the threshold a case of *uncertainty* about a matter of fact. What we have argued in this book is that our ambivalence about the threshold is not uncertainty, but *indecision*. There is no fact about the threshold to be uncertain about. If that is right, then we cannot hold that a possible world determines a value for the threshold.⁴⁴

Negotiated Contextualism

The view Justin Khoo calls "negotiated contextualism" comes closest to the view I am recommending here. The common ground, which Khoo calls the "conversational score," is modeled as a set of world/standard pairs, where a "standard" fixes the thresholds for gradable adjectives like "rich."⁴⁵ Unlike Silk, Khoo does not assume that a possible world fixes a standard, so his standards are nonfactual parameters. The "update

⁴³ In giving this sort of account, Silk has effectively given up on Assertoric Compatibility. He holds that in the cases of disagreement at issue, there is a single content (the diagonal proposition) that one party accepts and the other rejects. Given this, we do not need to say that the two parties are asserting compatible contents and disagreeing only in expressing competing normative views. At times, though, Silk does seem to endorse something like Assertoric Compatibility: "It cannot be a general requirement for discourse disagreement that the (intended) contents of the speakers' utterances be incompatible: the interlocutors in the PCS-discourses disagree even though they are talking about (or at least intending to talk about) different people" (Silk, *Discourse Contextualism*, 22).

⁴⁴ Silk sometimes says that speakers are "undecided" about what the abstract context should be, but, oddly, he seems to take this indecision to be compatible with the abstract context's being determined by a possible world (Silk, *Discourse Contextualism*, 132–33). Indecision as he thinks of it becomes a kind of uncertainty, to be distinguished from "ordinary factual uncertainty."

⁴⁵ Justin Khoo, "Quasi Indexicals," *Philosophy and Phenomenological Research* 100, no. 1 (2020): 13.

value” of an assertion is the set of world/standard pairs at which it is true, and the update proposed by an assertion is to intersect this set with the conversational score:

Thus, an assertion of “Jim is rich” proposes to make it settled in the context that Jim has enough personal wealth to count as rich in that context. It does so by proposing to eliminate all those world, wealth standard pairs from the score on which Jim’s wealth does not exceed the wealth standard.⁴⁶

Khoo’s account of the dynamics, then, matches what I am proposing, with some terminological differences: “conversational score” instead of “common ground,” “update value” instead of “content.” The use of sets of world/standard pairs is just what is needed to handle mixed cases without assuming (with Silk) that speakers take there to be a fact about the standard governing their conversation.⁴⁷

Despite this, Khoo regards his view not as a form of expressivism, but as a kind of contextualism: “negotiated contextualism.” He takes gradable adjectives (like “rich”) to be context-sensitive expressions—“quasi-indexicals”—because their contents “[depend] on a parameter of the context.”⁴⁸ I have argued, by contrast, that “rich” (or to be more precise “rich for an F ”) has a constant, context-independent content.

To understand this difference, we need to see what Khoo means by “content.” To avoid confusion, I will use the term “Khoo-content” when speaking of contents in Khoo’s sense. The *Khoo-content* of an expression at a context is defined as a set of possible-worlds intensions: For each standard g that appears in *some* world/standard pair from the context’s scoreboard, this set includes the intension that you get if you fix the standards parameter as g .⁴⁹ A Khoo-content is, essentially, a “cloud” of classical intensions. The Khoo-content of an expression like “rich” depends on context because different sets of standards are live options

⁴⁶ Khoo, “Quasi Indexicals,” 15.

⁴⁷ Khoo remarks: “For Silk, the crucial point is that the speakers have incompatible presuppositions about the relevant contextual parameter, while for me it is that the speakers propose incompatible updates to the context as part of their assertions (and so, for me the speakers may both agree about their presuppositions about how the context is—their disagreement is over how to change it)” (Khoo, “Quasi Indexicals,” 18).

⁴⁸ Khoo, “Quasi Indexicals,” 16.

⁴⁹ Khoo, “Quasi Indexicals,” 14.

in different contexts. As a result, different clouds of classical intensions will be associated with “rich” in different contexts. This is Khoo’s reason for classing his view as a form of contextualism.

However, it is only a good reason if these clouds of classical intensions are contents in more than just name. Traditionally, we derive the significance of an assertion from its assertoric content and facts about the context. In a Stalnakerian framework, for example, the significance of an assertion is its effect on the context set, which is completely determined by its content. The notion that plays this role in Khoo’s theory is not the Khoo-content, but the *update value* (a set of world/standard pairs). Khoo-content cannot play this role, because although an update value plus a context determines a Khoo-content, the converse is not true: A Khoo-content plus a context does not determine an update value, since it lacks information about which intension would be generated by which choice of standard. For example, if all you know is that an assertion has the Khoo-content

$$\{\{w_1, w_2\}, \{w_2, w_3\}\},$$

then you don’t know whether its update value is

$$\{(g_1, w_1), (g_1, w_2), (g_2, w_2), (g_2, w_3)\}$$

or

$$\{(g_2, w_1), (g_2, w_2), (g_1, w_2), (g_1, w_3)\}$$

or any number of others. So it is update values, not Khoo-contents, that can play the main theoretical role of assertoric content in Khoo’s theory.

What, then, is the role of Khoo-contents? Khoo suggests that they are, at least sometimes, the subjects of which truth and falsity are predicated. He has found, in several experiments, that respondents are hesitant to agree that in disputes like (42), at least one of the parties’ claims “has to be false.” We can explain this, Khoo thinks, by supposing that the respondents are evaluating the Khoo-content—a cloud of propositions—for falsity. They are reluctant to call it false because some of the propositions in the cloud are true.⁵⁰

Traditionally, assertoric content has had at least two theoretical roles: determining the significance of an assertion (the update) and being the

⁵⁰ Khoo, “Quasi Indexicals,” 15–16.

subject of which truth and falsity are predicated. It might seem tempting, then, to see Khoo as arguing that these roles must come apart, and introducing separate notions (update value and content, respectively) to play them. However, Khoo acknowledges that truth and falsity are often predicated of the *update value*, not the Khoo-content.⁵¹ For example, Caleb can join Ava's side in the metalinguistic negotiation of (42) by saying:

(44) What Ava says is true.

He is not saying that Ava's claim is true on all the standards that are in play in the scoreboard—presumably, he and Ava would both deny that. Rather, he is simply agreeing with Ava; he could equally well have said:

(45) Ava is right—Fred is rich.

Similarly, Delphine might summarize what is at issue in the conversation by saying:

(46) The issue is whether it's true that Fred is rich. Ava thinks it is, and Ben thinks it isn't.

Truth here is not being predicated of the Khoo-content: The issue is not whether it is *already settled* in the conversation that Fred counts as rich. Rather, these uses of "true" seem to operate transparently, in the sense that "It is true that Fred is rich" can be substituted for "Fred is rich" (and vice versa) without affecting the truth value the sentence. Delphine could equivalently have said:

(47) The issue is whether Fred is rich. Ava thinks he is, and Ben thinks he isn't.

Recognizing these points, Khoo holds that truth and falsity are sometimes predicated of the Khoo-content, and sometimes predicated of the update value.⁵² Because the semantic types of these objects are different, he reasons, there must be two distinct predicates "true"—one applying to clouds of propositions, and the other to update values.

⁵¹ Khoo, "Quasi Indexicals," 23–24.

⁵² Khoo, "Quasi Indexicals," 23–24.

Khoo thinks that the expressivist view, lacking this distinction, cannot explain respondents' hesitance to accept that at least one of the contending parties' claims in disputes like (42) is false.⁵³ However, it is possible to explain the data that motivates Khoo's theory without multiplying notions of content and truth predicates.

In Khoo's survey, the respondents were asked to what extent they agreed with the statement

(48) At least one of [Ava's or Ben's] claims has to be false.⁵⁴

The respondents are aware that questions about the truth of claims like "Fred is rich" can't be answered without making a decision about the threshold or standard of wealth that governs the conversation. At the same time, they are not *in* the conversation and thus not in a position to make the kind of bold conversational venture that Ava and Ben can make, with the aim of pushing the standard in their favored directions. So they might naturally reason as follows:

It's not for me to make decisions about their conversation. If I say that at least one of their claims has to be false, I might be asked which one. And even if I say that I don't know, that seems to imply that I think there's a fact of the matter *which* of them is false. So, I should not agree with the statement "at least one of their claims has to be false."

If we assume a transparent truth (and falsity) predicate, this reasoning is mistaken. A clear-thinking respondent would realize that the statement at issue is one whose truth is guaranteed *whatever* the standard may be, and hence one that can be accepted without making a decision about the standard. But the mistake is a subtle and understandable one. Avoiding it would require saying something sophisticated, like the following:

⁵³ Khoo, "Quasi Indexicals," 11n18. His target here John MacFarlane, "Vagueness as Indecision," *Aristotelian Society Supplementary Volume* 90, no. 1 (2016): 255–83. Khoo gives the same reason for resisting a Gibbardian interpretation of his update values ("Quasi Indexicals," 15n25).

⁵⁴ I have changed the names to align with my own example (42). Khoo worries about a potential scope ambiguity with the modal "has to" and repeats the experiment with "It must be that one of their claims is false," with no difference in the results ("Quasi Indexicals," 7). I am not sure why he didn't use the simpler statement "At least one of [Ava's or Ben's] claims is false."

- (49) At least one of their claims has to be false, but it is indeterminate which one is false.

It is not surprising that respondents without concepts like “indeterminate” would be tempted to reject (48).

Indeed, I suspect they will feel the same pressure to reject

- (50) It has to be the case either that (as Ava says) Fred is rich or that (as Ben says) he isn't,

or even

- (51) Either Fred is rich or he isn't.

These sentences do not use a truth predicate at all, so Khoo's explanation of respondents' reluctance to accept (48) will not generalize to (50) and (51). To explain these rejections without attributing any error, we would need more than a special truth predicate and a special notion of content; we would need a semantics that invalidates the classical law of excluded middle. Our error theory, on the other hand, does generalize: Speakers reject (50) or (51) because they confuse the question of *whether A or not-A* with the question of *whether it is settled that A or settled that not-A*.⁵⁵

To sum up, even if we spot Khoo his cloudy notion of “content” and the special truth predicate that goes with it, these will at best explain reluctance to accept (48); we will need to say something else about reluctance to accept (50) and (51). But any explanation of that will automatically give us an explanation of reluctance to accept (48), assuming a transparent truth predicate. I conclude, then, that we have no theoretical use for Khoo-contents or the special, nontransparent truth predicate that goes with them. And without these, there is no reason to think of Khoo's view as a form of contextualism, or “rich” as a quasi-indexical.

⁵⁵ Supervaluationists routinely appeal to a similar error theory to explain why speakers are sometimes reluctant to accept instances of excluded middle in borderline cases, given that these are tautologies on the supervaluationist account. The theory is that, when asked whether they accept “A or not-A,” speakers will sometimes answer as if they were asked about “definitely A or definitely not-A”—either because they do not clearly distinguish the two questions, as suggested here, or because, as Rosanna Keefe, *Theories of Vagueness* (Cambridge University Press, 2000), 164, proposes, they assume that they are not being asked about a tautology and reinterpret the question is a nontautological way.

It is worth noting one further difference between Khoo's negotiated contextualism and the expressivist theory defended here. On the expressivist view defended here, *all* lexical items are in principle plan-dependent. By contrast, Khoo regards quasi-indexicality as a feature of certain lexical items, paradigmatically gradable adjectives. Ordinary demonstratives and indexicals, he holds, are "directly" context-sensitive, and their behavior is not mediated by the scoreboard or subject to metalinguistic negotiation.⁵⁶ But there are certainly cases where the denotation of a demonstrative seems to be up for negotiation. Imagine a party of American men and women in 1900. One of the men says:

(52) In this country, we can all vote!

A woman replies:

(53) No, we cannot all vote!

Here, the issue is who counts as one of "us": The man may be planning to use "we" to apply just to the men present, and the woman is pushing back on this plan.⁵⁷ Metalinguistic negotiation does not seem to be restricted to any special class of expressions.

5.10 Conclusion

In making assertions and disagreeing with others' assertions, people often do express disagreement about how to use words. That is the important kernel of truth at the heart of theories of "metalinguistic negotiation." But in the literature, this kernel of truth is usually found mixed with some false chaff, to wit:

- The idea that, in disputes of this kind, the proposition asserted is not the locus of agreement.
- The idea that the two parties assert propositions that do not contradict each other.
- The idea that the disagreement is merely about language.

⁵⁶ Khoo, "Quasi Indexicals," 18.

⁵⁷ See Silk, *Discourse Contextualism*, 19 for a similar example.

These further commitments will seem inevitable so long as one makes certain orthodox assumptions about assertoric content. If Ava and Ben are asserting the kind of content that could be modeled as a set of possible worlds—a content whose truth or falsity is settled by how things are—then it is natural to suppose that “Fred is rich” expresses a different proposition when used by Ava than when used by Ben. That implies that the locus of disagreement is not the proposition that Fred is rich, but an expressed (but implicit) normative claim about how the word “rich” is to be used.

I have argued that the resulting position is untenable, and the only viable alternative is to drop the orthodox assumptions that make it inevitable and to model assertoric contents not as sets of worlds, but as sets of world/assignment pairs. Making that move allows us to see disagreements like (3) and (4) as what Plunkett and Sundell call “primary disagreements,” in which the locus of disagreement is the proposition asserted, while still seeing them as, in part, *practical* disagreements—disagreements in plan.

A great deal of theorizing about assertion starts from the idea that the primary point of assertion is to convey information. This is not to say that *every* assertion aims to convey information. Even in a context where everyone is presumed to know the result of last night's game, one might try to establish a topic of conversation or foster social bonding by asserting, "The Bears won big!" Similarly, a student's assertion that Sacramento is the capital of California is not intended to convey this information to the teacher, who is presumed to know it already; the point is rather to verify that the student knows this fact. Nonetheless, it is commonly held that the primary function of assertion is conveying information, and that even these noninformative exceptions are to be understood in terms of that primary role. I will call this the *informational model* of assertion. According to this model, the contents of assertions are *information contents*.

The main thrust of this book is that this picture must be rejected. Wherever vagueness and felicitous underspecification are in play—and that is just about everywhere—we cannot understand assertion as aiming to transfer information. There will not, in general, be a good answer to the question, "What information is the speaker intending to communicate?" The communicative effect of assertion must be understood as a hybrid one that aims not only to convey information, but to firm up plans for using words and concepts.

One might wonder, though, whether the informational model could be defended by observing that information can be probabilistic. The weather forecaster tells us that it is 60 percent likely to rain tomorrow; is this not information about the weather? Indeed, it is plausible that

the information we get from our own senses is probabilistic: A dimly glimpsed shape might be represented as 90 percent likely to be a leaf and 10 percent likely to be a lizard. We have seen that it is implausible that an assertion of “Bill is tall” is intended to communicate some specific worldly proposition—say, that Bill is greater than 190 cm in height. But one might still think that the assertion is meant to communicate information about Bill’s height. Just like the weather forecaster on the evening news, the speaker means to change the hearer’s probability distribution over possible heights for Bill in a predictable way, shifting the probability mass toward greater heights (see figure 6.1). Probabilistic information is still information, so this approach to vagueness is consistent with the informational model.




Figure 6.1: Shift in probabilities about Bill’s height in response to an assertion of “Bill is tall.”

In this chapter, we will examine two ways of working out this idea. Both use the Rational Speech Act (RSA) model of communication, a Bayesian pragmatic model that has been successful in explaining phenomena like scalar implicature, ambiguity, and reference resolution. I will argue that, despite its initial appeal, this approach has no plausible application to vagueness and indeterminacy. We cannot think of the intended effect of vague assertion as probabilistic information transfer. These arguments give further support to the expressivist alternative I have defended in this book.

6.1 Information and Communication

We must be careful at the outset to distinguish two questions. The first is whether the hearer gains any probabilistic information from the assertion. The second is whether any probabilistic information is *communicated*. These are very different questions.

To the first question, the answer is clearly yes. The hearer will go into the exchange with some prior probabilities about how tall Bill is, and the assertion will give the hearer a reason to update these probabilities, in something like the way shown in figure 6.1. So the hearer does get some information from the assertion. But that is not to say that information is *communicated*. After all, the hearer can get information from the sound made by a falling tree, and nobody would say that the tree has communicated anything. Similarly, when I say, “On va discuter la probabilité et la signification,” you may gain the information that I speak French with an American accent, but I haven’t *meant* or *communicated* this, even if I intended for you to notice it. Communication, as Paul Grice made clear, requires recognition of an intention that is, in part, an intention to be so recognized.¹

From the fact that one can gain probabilistic information from the vague utterance, then, nothing follows directly about what has been *meant* or *communicated*. To justify the claim that the *communicated* content is probabilistic, we need to show, not just that the listener’s credences change in response to the utterance, but that they change as a result of the listener’s recognition of the speaker’s intention that they change in this precise way.

What is needed, then, for a probabilistic account of vague meaning and communication, is an explanation of how the listener can come to recognize the credal change intended by the speaker. How can this be done?

6.2 The RSA Model

The RSA model seeks to answer this question. The model elaborates on a Gricean theme. On the traditional Gricean picture, the listener assumes that the speaker is trying to be cooperative and is therefore aiming to

¹ H. P. Grice, “Meaning,” *Philosophical Review* 66, no. 3 (1957): 377–88.

make the most informative utterance that is both true and relevant to the question.² Because the speaker's judgments about the truth and informativeness of utterances will depend on other assumptions the speaker is making, the listener can infer these assumptions from the speaker's choice of an utterance. In this way, the listener can recover the intended update, ruling out possibilities on which the utterance chosen would not have been the most informative one. This update is not essentially a probabilistic one, although once probabilities are eliminated, the probabilities of the uneliminated possibilities must be normalized.

The RSA model changes the listener's assumption. Instead of assuming that the speaker will always choose the most informative utterance, the listener assumes that the speaker will choose utterances with a probability that depends on their informativeness, selecting the most informative utterances most often, but sometimes selecting less informative ones.³ The soft-max choice rule is used to calculate the probability that utterance u will be used in situation s :

$$\Pr(u \mid s) \propto e^{\alpha \times I(u,s)}, \quad (6.1)$$

where $I(u,s)$ measures the informativeness of u in the event that s is true. Here, α is a parameter, sometimes called "temperature," that represents the degree to which the speaker's choice of utterance is determined by informativeness. When $\alpha = 0$, the choice of utterance is random and does not depend on informativeness at all. As α increases, the probability that the most informative utterance will be chosen increases. (In this chapter, we will assume that $\alpha = 4$.)⁴

Because there is always a chance the speaker will select a nonoptimal utterance, the listener cannot react to an utterance by ruling out possibilities on which the utterance is not the most informative one the speaker could have chosen. Thus, we get an update of the listener's credences that is not entirely attributable to the elimination of any live possibilities. But

² We'll ignore considerations of economy and manner for now.

³ This is a simplification; the RSA model also takes into account the cost of an utterance in determining its overall utility. To keep things simple, I'll be working with sets of utterances that do not differ in cost, so we can ignore cost and just focus on informativeness.

⁴ This is the value used in Daniel Lassiter and Noah D. Goodman, "Adjectival Vagueness in a Bayesian Model of Interpretation," *Synthese* 194, no. 10 (2017): 3801–36.

even though the intended effect on the listener's mentality is now essentially probabilistic, we still have the basic structure of Gricean reasoning. The listener assumes the speaker is cooperative, and thus is more likely to choose an utterance u the more informative this utterance would be. Based on this assumption, the listener estimates the probability that the speaker will utter u , conditional on each possible state of affairs s , $\Pr(u | s)$. This information allows the listener to derive new probabilities for states of affairs by exploiting Bayes' theorem:

$$\Pr'(s | u) = \frac{\Pr(u | s) \times \Pr(s)}{\Pr(u)}, \quad (6.2)$$

where

$$\Pr(u) = \sum_{s \in S} \Pr(u | s),$$

S being the set of possible situations. So, the probability the hearer assigns to s , on learning that u has been uttered, will go up if u is more likely to have been uttered by a cooperative speaker conditional on s than conditional on a tautology, and it will go down if it is less likely.

It remains only to describe how we calculate the informativeness of utterances, $I(u, s)$:

$$I(u, s) = \ln \Pr(s | u \text{ is true}). \quad (6.3)$$

The natural log ($\ln$) of a probability can vary from $-\infty$ ($= \ln 0$) to 0 ($= \ln 1$). Suppose we are concerned with the roll of a six-sided die and the prior probability of each side coming up is $1/6$. In fact, the die came up 3. Call that state of affairs s . A maximally informative utterance u ($I(u, s) = \ln 1 = 0$) would then be one whose truth makes s certain: "The die landed on 3." A less informative utterance would be "The die landed on an odd number." The probability of the die having landed on 3 conditional on the truth of this utterance is $1/3$, so $I(u, s) = \ln 1/3 \approx -1.099$. An even less informative utterance would be "The die landed on a number between 1 and 6." Here, $I(u, s) = \ln 1/6 \approx -1.792$.

What would be a minimally informative utterance? One might think it would be a tautology, but the probability of the die landing on 3, conditional on a tautology, is $1/6$, so a tautology has an informativeness of approximately -1.792 . A minimally informative utterance ($I(u, s) = -\infty$) would be one whose truth is *incompatible* with the actual state of affairs,

such as “The die landed on 4” or “The die landed on 20.” Indeed, if $\Pr(s)$ is already very high, a tautology can have a very high informativeness score. It might seem more sensible to measure informativeness by the *degree of increase* in credence in s induced by conditionalizing on the truth of u , so tautologies all had minimal informativeness. However, if tautologies always had minimal informativeness, then uttering a tautology (or, equivalently, remaining silent) could never be preferable to uttering a patent falsehood.

The basic idea, then, is that an utterance is more informative to the extent that conditionalizing on its truth would give the listener higher credence in the *actual* answer to the question under discussion (QUD), as known by the speaker. So, if the QUD is “What is for dinner?” and the answer (as known by the speaker) is beef, then “beef” is a more informative utterance than “meat.”

What if the speaker doesn’t know what is for dinner, but only has probabilities? We can complicate the model to handle this case by defining informativeness in terms of how closely the listener’s probabilities, conditional on the truth of the utterance, accord with the speaker’s, as measured by the Kullback-Leibler divergence. But the simpler model will suffice for us here, and the complex model reduces to the simpler one in the case where the speaker has full knowledge of the relevant facts.

We are now in a position to see how reasoning probabilistically about the speaker’s utterance, on the assumption that the likelihood of the speaker’s choosing any utterance is proportional to its informativeness, can give us a probabilistic update that goes beyond simply ruling out the possibilities that are incompatible with the truth of the utterance.

Suppose the listener’s priors about the menu are evenly divided among three options, beef, chicken, and ratatouille, each with probability $1/3$. And let’s assume, somewhat artificially, that there are three possible utterances available to the speaker: “meat,” “beef,” and “food.” Simply conditionalizing on the truth of the utterance “meat” would eliminate ratatouille and give probability $1/2$ each to beef and chicken. But for the pragmatic listener who reasons about the speaker’s intentions and assumes the speaker to be cooperative, hearing the speaker say “meat” will actually *reduce* the probability of beef, because “meat” would be a much less informative utterance than “beef” in that case. Our expectation, then, is that the speaker is very unlikely to utter “meat” if beef is for

dinner. Hearing the speaker utter “meat” should make us very confident that we will have chicken.

Let’s run the numbers. First, we calculate the informativeness of each utterance, given each state of affairs, using equation (6.3):

$$I(\text{“meat”}, \text{beef}) = \ln 1/2.$$

$$I(\text{“beef”}, \text{beef}) = \ln 1.$$

$$I(\text{“food”}, \text{beef}) = \ln 1/3.$$

$$I(\text{“meat”}, \text{chicken}) = \ln 1/2.$$

$$I(\text{“beef”}, \text{chicken}) = \ln 0.$$

$$I(\text{“food”}, \text{chicken}) = \ln 1/3.$$

$$I(\text{“meat”}, \text{ratatouille}) = \ln 0.$$

$$I(\text{“beef”}, \text{ratatouille}) = \ln 0.$$

$$I(\text{“food”}, \text{ratatouille}) = \ln 1/3.$$

Given our assumption that the speaker will prefer more informative utterances (equation (6.1) with $\alpha = 4$), we have

$$\Pr(\text{“meat”} \mid \text{beef}) \approx 0.06.$$

$$\Pr(\text{“beef”} \mid \text{beef}) \approx 0.93.$$

$$\Pr(\text{“food”} \mid \text{beef}) \approx 0.01.$$

$$\Pr(\text{“meat”} \mid \text{chicken}) \approx 0.84.$$

$$\Pr(\text{“beef”} \mid \text{chicken}) = 0.$$

$$\Pr(\text{“food”} \mid \text{chicken}) \approx 0.16.$$

$$\Pr(\text{“meat”} \mid \text{ratatouille}) = 0.$$

$$\Pr(\text{“beef”} \mid \text{ratatouille}) = 0.$$

$$\Pr(\text{“food”} \mid \text{ratatouille}) = 1.$$

From this, we can determine absolute probabilities for the different possible utterances:

$$\Pr(\text{“meat”}) = (0.06/3) + (0.84/3) + (0/3) \approx 0.30.$$

$$\Pr(\text{“beef”}) = (0.93/3) + (0/3) + (0/3) \approx 0.31.$$

$$\Pr(\text{“food”}) = (0.01/3) + (0.16/3) + (1/3) \approx 0.39.$$

Now we can plug these values into equation (6.2) to get the probability of a state of affairs conditional on each utterance:

$$\begin{aligned}
 & \Pr(\text{beef} \mid \text{"beef"}) \\
 &= \frac{\Pr(\text{"beef"} \mid \text{beef}) \times \Pr(\text{beef})}{\Pr(\text{"beef"})} \\
 &\approx \frac{0.93 \times 0.33}{0.31} \\
 &\approx 0.99 \\
 & \Pr(\text{beef} \mid \text{"meat"}) \\
 &= \frac{\Pr(\text{"meat"} \mid \text{beef}) \times \Pr(\text{beef})}{\Pr(\text{"meat"})} \\
 &\approx \frac{0.06 \times 0.33}{0.30} \\
 &\approx 0.07.
 \end{aligned}$$

So, even though the literal content of "meat" is perfectly compatible with beef, an utterance of "meat" by a speaker assumed to be cooperative dramatically reduces the probability of beef.

That's the basic idea of the RSA framework. If the listener's priors are common knowledge between speaker and listener, then the speaker can expect that the listener will adjust probabilities this way, so the speaker can rationally intend that the listener recognize her intention for them to change in this way.⁵

This model hews closely to the Gricean picture, while substituting a probabilistic update for the standard nonprobabilistic one. It has been applied to standard pragmatic phenomena such as scalar implicatures and ambiguity and reference resolution, with good results, in the sense that the model's predictions (with suitable values for parameters) agree well with experimentally tested responses. But things get more complex when we try to apply the RSA model to vague utterances, like "Bill is tall."

⁵ You might wonder why informativeness is defined by the effect of the utterance on a "literal listener," who simply conditionalizes on the truth of an utterance. Shouldn't the speaker be thinking about pragmatic dimensions of informativeness as well? That can be captured by iterating the model, considering a level-2 pragmatic speaker who is evaluating informativeness to our level-1 pragmatic listener, and a level-2 pragmatic listener who is drawing inferences from the level-2 pragmatic speaker's choice of utterance. And we can iterate to as many levels as we like. In the cases we will discuss here, iterating does not make much difference, so I stay on the first level.

6.3 Applying RSA to Vague Language

It is not implausible that we reach for vague sentences when we want to induce a change in the hearer's probabilities, without committing ourselves to any definite claim about the world. Consider this example, due to Paul Égré, Benjamin Spector, Adèle Mortier, and Steven Verheyen (henceforth ESMV).⁶ When asked to report how many people attended a party, you could respond in two very different ways:

- (1) Around 70 people were at the party.
- (2) Between 60 and 80 people were at the party.

Asserting (2) makes a definite claim about the world; it proposes ruling out all possibilities where there were fewer than 60 or more than 80 people at the party. Accepting this assertion will affect the hearer's credences, by forcing all the probability mass between 60 and 80, but a large variety of credence distributions would be compatible with accepting (2). For example, the hearer could suppose that all numbers of guests between 60 and 80 are equiprobable, but many other distributions would be compatible with the speaker's intentions. So the intended effect of asserting (2) does not seem to be probabilistic. Any probabilistic effect is just the by-product of a nonprobabilistic one—the ruling out of alternatives.

In asserting (1), by contrast, a speaker plausibly intends the hearer to have a certain kind of distribution of credences, centered on 70 and dropping off above and below, but not artificially bounded by 60 or 80 (see figure 6.2). The intended effect of asserting (1) is not, then, just ruling out possibilities; rather, it seems intrinsically probabilistic. In asserting (1), one intends to shape the hearers' credences in a particular way.

To reiterate the main point of section 6.1, the question here is not whether hearers' credences will reliably change in a certain way, but whether the speaker's meaning—a communicative effect the speaker intends the hearer to grasp—can consist in such a change. For that to be the case, it must be possible for the hearer to work out the probabilistic change that the speaker intends. The RSA model looks promising as

⁶ Paul Égré, Benjamin Spector, Adèle Mortier, and Steven Verheyen, "On the Optimality of Vagueness: 'Around,' 'Between' and the Gricean Maxims," *Linguistics and Philosophy* 46, no. 5 (2023): 1075–130.




Figure 6.2: Intended effect on credences of asserting “Around 70 people were at the party.”

an idealized account of how that might be done. If we can quantify the informativeness of various candidate utterances and assume that the probability that the speaker chooses an utterance is proportional to its informativeness, then we can derive posterior probabilities in the way described in section 6.2.

Recall that when the question under discussion is Bill’s height, an utterance is informative to the extent that conditionalizing on its truth gives the listener high confidence in Bill’s actual height (which we assume to be known by the speaker). So if the speaker knows that Bill is 200 cm tall, while the listener has a normal distribution over possible heights, then the utterance “Bill is over 190 cm tall” is more informative than “Bill is over 180 cm tall,” and “Bill is exactly 200 cm tall” is maximally informative.

But how do we measure informativeness for the vague utterance “Bill is tall”? How informative an utterance like “Bill is tall” is would seem to depend not just on Bill’s height, but on how tall one has to be to count as “tall” (in this context), and this is not known by the listener or the speaker.

Here, we have some choices. I’m going to start by sketching the approach of Daniel Lassiter and Noah D. Goodman.⁷ Later, in response to an objection, I’ll consider ESMV’s different way of tackling the problem.

⁷ Lassiter and Goodman, “Adjectival Vagueness in a Bayesian Model of Interpretation.”

Assume a standard degree semantics, according to which that a vague utterance like “Bill is tall” is true just in case Bill’s height exceeds the threshold θ for counting as tall (see chapter 4, section 4.2). We can then measure informativeness as before, except that it is now relative not just to Bill’s height, but also to the threshold:

$$\begin{aligned} I(u, h = x, \theta = t) &= \ln \Pr(h = x \mid u \text{ is true} \ \& \ \theta = t) \\ &= \ln \Pr(h = x \mid h \geq t), \end{aligned}$$

where h is Bill’s height and θ is the threshold for counting as tall. “Bill is tall,” construed on the degree semantics model, will be more informative the closer θ is to h , so long as θ is lower than h . (Once it becomes higher, the utterance becomes false and the informativeness score goes to $\ln 0 = -\infty$.)

Table 6.1 summarizes some informativeness calculations with different parameters, along with the choice probabilities they induce, given the utterance alternatives “Bill is tall,” “Bill is not tall,” and a tautology (= saying nothing).⁸

What we have so far is a way for a pragmatic listener to calculate $\Pr(u \mid h = x \ \& \ \theta = t)$ for various values of x and t . But how is the listener, who knows neither h nor θ , supposed to use these probabilities to figure out what the speaker is trying to convey about Bill’s height? As before, we appeal to Bayes’ theorem, but now we have not just a state of the world (Bill’s height), but also a threshold to consider:

$$\begin{aligned} \Pr'(h = x \ \& \ \theta = t \mid u) \\ = \frac{\Pr(u \mid h = x \ \& \ \theta = t) \times \Pr(h = x \ \& \ \theta = t)}{\Pr(u)}. \end{aligned} \tag{6.4}$$

Equation (6.4) shows how the a listener with a prior probability distribution over height/threshold pairs can infer a posterior probability

⁸ Note that increasing θ while keeping h fixed increases the probability of “Bill is tall” relative to a tautology, because the informativeness of “Bill is tall” is sensitive to θ , while that of a tautology is not. By contrast, increasing h while keeping θ fixed does not change the probabilities of “Bill is tall” and a tautology, because it decreases the informativeness of both in a correlated way. Arguably, this is counterintuitive. Cyang Qing and Michael Franke, “Gradable Adjectives, Vagueness, and Optimal Language Use: A Speaker-Oriented Model,” *Proceedings of SALT 24* (2014): 35, takes this to be a serious objection against the Lassiter and Goodman model, arguing that empirically, “the greater the degree [of height], the more likely that the positive form will be used.”

Table 6.1: Sample informativeness calculations and the induced utterance-choice probabilities (given the alternatives “Bill is tall,” “Bill is not tall,” and a tautology), with $\alpha = 4.0$. Numbers are rounded to two decimal places.

u	x	t	$I(u, h = x, \theta = t)$	$Pr(u \mid h = x \ \& \ \theta = t)$
Bill is tall	190	185	-2.42	0.93
Bill is not tall			$-\infty$	0.00
Tautology			-3.05	0.07
Bill is tall	190	190	-1.55	1.00
Bill is not tall			$-\infty$	0.00
Tautology			-3.05	0.00
Bill is tall	200	185	-5.20	0.93
Bill is not tall			$-\infty$	0.00
Tautology			-5.83	0.07
Bill is tall	200	190	-4.33	1.00
Bill is not tall			$-\infty$	0.00
Tautology			-5.83	0.00
Bill is tall	200	205	$-\infty$	0.00
Bill is not tall			-5.83	0.50
Tautology			-5.83	0.50

distribution from the speaker’s utterance, on the assumption that the speaker will produce utterances in proportion to the exponential of their informativeness. (We will return later to the assumption that the listener has a prior probability distribution over height/threshold pairs.)

Figure 6.3 shows how an utterance of “Bill is tall,” “Bill is not tall,” or a tautology affects the pragmatic listener’s credences about Bill’s height and the value of the threshold θ . As one can see, even a tautology can lead to changes in credence, since the speaker’s choice to utter a tautology (or nothing) rather than “Bill is tall” or “Bill is not tall” tells us something about Bill’s height.




Figure 6.3: Effect of uttering (a) "Bill is tall," (b) "Bill is not tall," or (c) a tautology, on Lassiter and Goodman's model, where the speaker's alternatives are "Bill is tall," "Bill is not tall," and a tautology, and $\alpha = 4.0$.

6.4 A Conceptual Objection

Although it gives plausible predictions about how credences can be predicted to change in response to vague utterances, there is a serious conceptual problem with this account.

As we have defined it, informativeness depends on θ . If θ were known to the speaker, then, on the assumption that the speaker is more likely to produce more informative utterances than less informative ones, we could infer something about θ from the speaker's choice of utterance. But if—as seems to be the case— θ is unknown to the speaker, then the

speaker's choice of utterance can hardly be evidence about the value of θ . If the informativeness of the utterance depends on θ , the speaker's choice of utterance cannot depend on its informativeness, because the speaker does not know the value of θ .

To put it another way, if we define informativeness the way Lassiter and Goodman do, then speakers aren't going to know how informative different utterances would be, so they *can't* choose utterances in proportion to their informativeness.

Here is how Lassiter and Goodman describe the reasoning process of the "pragmatic listener":

Our pragmatic listener will now instantiate assignments of values V from a prior distribution $P_{L_1}(V)$ and thread them through the iterated reasoning procedure, considering how likely it is that the speaker would have produced the utterance that he did, *if the variables were resolved as they are in V* .⁹ [Here, V is playing the role of our θ .]

But this is not the right counterfactual. If the speaker is not epistemically tracking the value of θ , then counterfactual variations in θ would have made no difference to the speaker's choice of utterance.

In their introduction to the RSA model, Scontras, Tessler, and Franke describe the pragmatic reasoning process in a similarly incoherent way:

In other words, L_1 simulates S_1 's behavior for the possible combinations of states and thresholds, then weights states and thresholds in proportion to the probability that they would lead S_1 to choose the utterance that was encountered.¹⁰

But if the speaker (S_1) does not know the threshold, then the speaker's behavior can't be sensitive to changes in the threshold, as this account seems to assume. Qing and Franke acknowledge that the RSA model "assumes that the speaker knows the exact threshold for the positive form, which is arguably [!] not the case in reality."¹¹

⁹ Lassiter and Goodman, "Adjectival Vagueness in a Bayesian Model of Interpretation," 3817, emphasis in the original.

¹⁰ Gregory Scontras, Michael Henry Tessler, and Michael Franke, "A Practical Introduction to the Rational Speech Act Modeling Framework" (2021), 12, <https://doi.org/https://doi.org/10.48550/arXiv.2105.09867>.

¹¹ Qing and Franke, "Gradable Adjectives, Vagueness, and Optimal Language Use," 35.

Lassiter and Goodman could resort to saying that the RSA model is just a mathematical model and doesn't need to "make sense" as a piece of argumentation. But to say this would be to give up the Gricean motivation for the project. We want to describe a way in which listeners could come to *recognize* how a speaker intends for them to modify their credences. Recognition is a kind of knowledge and cannot be produced by nonsensical reasoning.

6.5 Fixing the Bug

To fix this bug, we need a measure of informativeness that is epistemically accessible to the speaker.

One solution would be to define informativeness relative not to a threshold and height, but to a probability distribution over threshold/height pairs. (This solution would also handle cases where the speaker is ignorant of Bill's exact height.) To do this, we can use the Kullback-Leibler divergence, a standard measure of the degree to which one probability distribution matches another. The informativeness of an utterance would depend on the degree to which the listener's posterior distribution over height/threshold pairs, after conditionalizing on the truth of the speaker's utterances, matches the speaker's priors.

Another approach (used in ESMV) is to have the literal listener generate a distribution over heights by conditionalizing the height/threshold pair distribution on the truth of the utterance, and then summing the probabilities of pairs with the same height component to get a distribution over heights rather than pairs. For example, if there are n possible threshold values, $t_1 \dots t_n$, then we calculate

$$\Pr(h = x \mid u \text{ is true}) = \sum_{i=1}^n \Pr(h = x, \theta = t_i \mid u \text{ is true}).$$

We can then define informativeness without reference to θ , in the same way we did in equation (6.3), as

$$I(u, h = x) = \Pr(h = x \mid u \text{ is true}).$$

Figure 6.4 shows what happens on the ESMV model with assertions of "Bill is tall," "Bill is not tall," and a tautology. Note that here, the threshold priors and posteriors overlap because the pragmatic listener is not inferring anything about the threshold.




Figure 6.4: Effect of uttering (a) "Bill is tall," (b) "Bill is not tall," or (c) a tautology, on the ESMV model, where the speaker's alternatives are "Bill is tall," "Bill is not tall," and a tautology, and $\alpha = 4.0$.

6.6 Further Problems for RSA Models

Either of these fixes would avoid the conceptual problem with Lassiter and Goodman's view, while keeping many of its good points. However, there are further problems that affect these approaches, as well as Lassiter and Goodman's.

Assertion Is Not Idempotent

The first problem is that, on these views, asserting "Bill is tall" is not idempotent. Each time the speaker asserts "Bill is tall," the pragmatic

listener's probability curve for Bill's height moves up a bit. Figure 6.5 illustrates this phenomenon for both the Lassiter/Goodman model and the ESMV model.¹² The first 10 iterations are shown. Subsequent assertions drive the height posteriors farther to the right, without converging to a fixed point.




Figure 6.5: Effect of repeated assertions of “Bill is tall” on Lassiter and Goodman’s model (a) and the ESMV model (b), with $\alpha = 4.0$. Darker lines indicates posteriors for later assertions, lighter lines for earlier assertions.

The reason is obvious: The intended posterior height and θ distributions depend on the prior height and θ distributions. Each assertion changes these, so the hearer’s estimate of height gradually increases.

Is failure of idempotence necessarily a bad thing? Some utterances—for example, “I put another uranium brick on the pile” or “I have said nothing so far”—are clearly not idempotent. But “Bill is tall” does not seem to have the special features that account for these failures of idempotence. Saying “Bill is tall,” having one’s assertion accepted, and then saying “Bill is tall” again would be a very odd thing to do. (One would naturally reply: “Yes, we know, you already said that.”) But on the accounts

¹² Note that the Lassiter/Goodman model assumes that h and θ are probabilistically independent for the listener. We cannot make this assumption if we hope to model iterated assertions, because even if they are independent at the beginning, they will not be independent after an assertion has been made. The fix is just to let the listener’s priors and posteriors be defined over (height, threshold) pairs.

we are considering, it is predicted to be a perfectly normal, cooperative speech act. If the repeated assertions are accepted, the result will be a steady increase in the listener's estimate of Bill's height. The ESMV account even predicts that repeated assertions of a tautology could result in a steady *decrease* in the listener's estimate of Bill's height. These are very odd predictions, and I think they count against these models.

Unrealistic Assumption of Common Ground

These models assume that the listener's priors (over both heights and thresholds) are common knowledge between the speaker and the listener. Otherwise, the speaker and hearer might (for all they know) have very different conceptions of the informativeness of an utterance in a state of affairs, and the listener would not be able to infer anything from the speaker's choice of utterance.

This assumption seems unrealistic. In many normal settings where vague language is used, the speaker doesn't know the listener's priors. For example, consider a guide in the Sydney Zoo who is showing the Australian section to visitors from all around the world. The guide may have no idea what the listener's distribution of credences over the heights of wallabies looks like. If the RSA model were correct, this should be a situation in which the guide cannot expect her assertion of

(3) That is a tall wallaby, but this one here is short.

or

(4) Most of the tall wallabies are in the other part of the park.

to be understood. The listeners will be able to work out the intended credal update only if they take the speaker to know what their priors are. In fact, however, lack of common knowledge about these credences is no obstacle to understanding or asserting (3). Such an assertion can help listeners calibrate their use of "tall" and "short," as applied to wallabies, and form their expectations about likely wallaby heights.

Informativeness Is Overemphasized

Both models predict that a speaker with full knowledge of Bill's height will *almost never* choose "Bill is tall" when more precise utterances like

"Bill is n cm tall" are available. On the Lassiter/Goodman model, for example, the probability that the pragmatic speaker will utter "Bill is 190 cm tall" instead of "Bill is tall" or a tautology, when the speaker knows that Bill is 190 cm tall and the threshold for counting as "tall" is 180 cm, is approximately 0.99998.¹³ This nearly invariable preference for exact assertions makes sense on an informational model of assertion. If speakers are trying to be maximally informative about Bill's height, they should choose "Bill is tall" over "Bill is 190 cm tall" only when their knowledge of Bill's height is inexact.

In fact, though, even speakers with exact knowledge of Bill's height will often find reasons to assert "Bill is tall" instead of "Bill is 190 cm tall." Chris Barker's well-known example of the use of "Feynman is tall" to fix the usage of "tall" is one such case.¹⁴ Another would be a situation in which we have agreed to select only tall people for the department basketball team. A realistic account of the pragmatics of vague expressions should be able to handle such cases.

One possible fix would be to define informativeness in terms of the Kullback-Leibler divergence of probability distributions over height/threshold pairs. That would be a way of formalizing the idea that utterances like "Feynman is tall" serve to convey metalinguistic information (here, about the threshold for "tall"), as well as information about height. But I think this is the wrong fix, for reasons I'll get to next.

Priors Over Thresholds

Both models presuppose a prior probability distribution over thresholds. (They do different things with it, but they both require it.) However, subjective probability distributions over thresholds make sense only for parties who assume that there is a fact about where the threshold lies. Only in that case does it make sense to ask about the probability that

¹³ As observed in n3, we are assuming for simplicity that all utterances are equally costly. However, even if we take "Bill is 190 cm tall" to be three times as costly as "Bill is tall," the Lassiter/Goodman model still predicts a 0.95 probability that the speaker will utter it instead of "Bill is tall." Its high informativeness outweighs its extra syllables. In any case, the point being made here could also be made about pairs of utterances that seem equally costly, like "Bill is pretty tall" and "Bill is six feet tall."

¹⁴ Chris Barker, "The Dynamics of Vagueness," *Linguistics and Philosophy* 25, no. 1 (2002): 1–36.

the threshold is, say, greater than 180 cm. Where we don't think there is a fact, subjective credences don't make sense.

In general, speakers do not think there is a fact about where the threshold lies. This is reflected in their unwillingness to express their ambivalence in borderline cases using ordinary locutions of subjective uncertainty.¹⁵ When we take Bill to be a borderline case of a tall man, we don't say

- (5) It's possible that he's tall, or
- (6) He's probably tall, but there's a chance he isn't.

Instead, we express our ambivalence using the language of indecision:

- (7) We could count him as tall.
- (8) We should probably count him as tall.

If we do not have credences about the value of θ , then it is hard to see how RSA models could accurately represent our pragmatic reasoning. Lassiter and Goodman claim not to be committed to the view that there is a fact about the threshold governing a context; they intend to be giving a model of the psychology of interpretation, not a metaphysical view.¹⁶ Fair enough, but the interest of the RSA model is that it purports to give a rational path by which speakers and listeners can coordinate on an intended probabilistic update. The reasoning involved requires the parties to assume that priors are defined for the threshold. Leaving aside issues of metaphysics, this is not a psychologically plausible assumption.

6.7 An Alternative

In chapter 3, I argued that we should not think of the threshold θ as fixed by any aspect of the context; at best, context constrains it to a *range*. There is no fact about which value in that range is the threshold, and it makes no sense to have subjective credences about the location of the threshold. The range represents not our uncertainty, but our indecision about where to put the threshold—indecision we can reduce not by discovery, but only by decision and further stipulation. Vague

¹⁵ See chapter 3, section 3.1.

¹⁶ Lassiter and Goodman, "Adjectival Vagueness in a Bayesian Model of Interpretation," 22–23.

utterances express hybrid states that combine doxastic and deliberative aspects. The intended effect of a vague assertion is to update a common ground that can be represented as a set of world/threshold pairs (or, more generally, a set of world/assignment pairs). In this framework, the intended effect of asserting “Bill is tall” is to remove from the context set any world/threshold pairs (w, θ) such that Bill’s height in w does not meet or exceed θ .

This model avoids all the problems faced by the RSA account. First, it secures idempotence: Asserting “Bill is tall” twice in a row produces the same update to the common ground as asserting it once. (Namely, all world/threshold pairs that would make this sentence false are rejected.)

Second, it does not make unrealistic assumptions about the common ground. In particular, it does not assume that the listener’s priors about Bill’s height are common ground.

Third, it allows us to make sense of both informational and noninformational updates. Asserting “Bill is tall” when it is common ground that the threshold for being tall is 190 cm, for example, will be an informational update: It will tell us that Bill is at least 190 cm in height. By contrast, in a situation where the threshold is indeterminate (but constrained to the range 180–200 cm), and it is common knowledge that Bill is 195 cm tall, asserting “Bill is tall” will have a noninformational effect—that of constraining the threshold to the range 180–195 cm. Accepting the assertion in that situation will amount to adopting a joint plan to use “tall” in a certain way. In the most common kind of case, where the common ground pins down neither the threshold nor Bill’s height, we have an update that is neither purely informational nor purely plannish. The update doesn’t rule out a threshold of 195 cm, and it doesn’t rule out Bill’s being 194 cm tall, but it does rule out the *combination* of a threshold of 195 cm and a height of 194 cm. This update is neither purely informational nor purely plannish.

Finally, our model avoids assuming that speakers or listeners have priors over thresholds (let alone *shared* priors). Indeed, our working assumption is that speakers and listeners do not have priors over thresholds, because they do not think there is a fact about where the thresholds lie.

With this alternative in mind, it is worth pausing to recall why the RSA model seemed like a tempting answer to the problem about vagueness and communication. If we start with an informational model of assertion, the puzzle about vagueness and communication will reduce to this question: What information can a speaker hope to impart to a listener with a vague assertion? In chapter 1, we argued that there is no way for speaker and hearer to coordinate on a particular information content—a possible-worlds intension. But it does not seem quite right to say that no information about the world is communicated. Those who hear a credible source assert “Bill is tall” will increase their credence that Bill is, say, greater than 200 cm. They won’t be able to rule out any specific height, but their credences will change, in a way that can be predicted and intended by the speaker. The thought, then, is that we can keep the idea that an assertion proposes an informational update if we allow the update to be probabilistic.

In addition, vague sentences seem to convey a distinctive kind of probabilistic information.¹⁷ Speakers will typically use vague sentences like

(9) There are around 100 jelly beans in the jar

only when their probabilities peak around 100 jelly beans and drop off gradually in both directions. It would be quite odd to assert (9) if you knew only that the number of jelly beans in the jar was determined by a random process, with equal chances for each number of jelly beans between 80 and 120. In that case, you’d say, “There are between 80 and 120 jelly beans in the jar.” It would also be odd to assert (9) if you had 0.5 credence that there were 100 jelly beans, 0.25 that there were 80, and 0.25 that there were 120. So there is *something* to the idea that, in many cases, probabilistic information is conveyed by vague utterances, and is intended to be conveyed. If we hold on to the informational model of assertion, on which the content of an assertion is the information it seeks to get across, then our only alternative will be to identify this probabilistic information as the content of the assertion.

The problem is that, absent the problematically strong assumptions of the RSA model, we don’t have anything like a recipe that listeners

¹⁷ Égré et al., “On the Optimality of Vagueness.”

can use to recognize the specific probabilistic update intended by the speaker. So we cannot think of this change as the intended illocutionary effect of the assertion. The expressivist account, by contrast, offers us an account of the intended update that the speaker and hearer are in a position to coordinate on. It's just not an *informational* update, and it's entirely nonprobabilistic.

Does this leave us with no explanation of the probabilistic effects of vague utterances? I don't believe it does: The expressivist has resources here. This will be easiest to see if we consider a particular case. Suppose that the common ground rules out setting a threshold for "tall" below 185 cm or above 200 cm, but it does not rule out any threshold between 185 and 200 cm. Now consider two possible assertions:

(10) Bill is tall.

(11) Bill is at least 185 cm tall.

Intuitively, even though (10) doesn't rule out any possible height for Bill that (11) doesn't rule out, these assertions have different significances. But in what does their difference consist?

On the RSA account, the two assertions carry different information. To keep things simple, let's assume that the relevant alternatives to these assertions include their negations and a tautology, and the listener's credences about Bill's height are commonly known and cluster in a normal distribution centered on 185 cm. Then the RSA account (in Lassiter and Goodman's version) predicts that if (11) is asserted, pragmatic listeners should put most of their credence on Bill being between 185 and 188 cm, while if (10) is asserted, they should put most of their credence on Bill being more than 188 cm tall (see figure 6.6).

Intuitively, this seems right. In a situation where we're sure that the threshold for "tall" is not less than 185 cm, we'll think Bill is taller if he is said to be "tall" than if he is said to be at least 185 cm. It is a strength of the RSA account that it vindicates this intuitive judgment. Can the expressivist account do the same?

On the expressivist account, both assertions—(10) and (11)—rule out the possibility that Bill is less than 185 cm tall (which, recall, is already accepted as the lower bound for the threshold for "tall"). And that is *all* that (11) does. But (10) does something further. It amounts to a commitment to count anyone who is at least as tall as Bill as "tall,"




Figure 6.6: Effect of uttering (a) "Bill is tall" or (b) "Bill is at least 185 cm tall," on the ESMV model, where the speaker's alternatives are "Bill is tall," "Bill is at least 185 cm tall," "Bill is not tall," "Bill is not at least 185 cm tall," and a tautology, and $\alpha = 4.0$.

no matter how tall Bill turns out to be. For example, it precludes later discovering that Bill is 195 cm tall and setting the threshold for "tall" at 200 cm. By contrast, (11) does not preclude this.

Under what conditions would this additional commitment be rational to undertake? It will be rational if the speaker is confident that it will not prove useful, in the future development of the conversation, to put the threshold for "tall" above Bill's height. The most obvious way the speaker can be confident of this is by having a very high estimation of Bill's height. So it is quite natural for a listener to infer from the speaker's decision to assert (10) instead of (11) that the speaker thinks that Bill is *significantly* taller than 185 cm. By contrast, if the speaker asserts (11), then it is natural to simply conditionalize on this information. This explains why a listener might estimate Bill to be taller after an assertion of (10) than after an assertion of (11).

However, one can imagine circumstances in which it is rational for a speaker to undertake the additional commitment in (10) even while thinking that Bill is not significantly more than 185 cm in height. For example, suppose that the topic of this conversation is who is to be selected for the departmental basketball team, that it is agreed that only tall people will be selected, and that it is commonly known that

the speaker has a strong interest in having Bill on the team, regardless of Bill's height. Then the inference that the speaker must have a high estimate of Bill's height is blocked. In this situation, the expressivist account does not predict the difference between "Bill is tall" and "Bill is at least 185 cm tall" depicted in figure 6.6. And I think that is a correct prediction.

Thus, not only can the expressivist account explain some of the empirical data that motivate the RSA account; it can also explain variations in that data that the RSA account cannot explain.

Bibliography

- Austin, J. L. *How to Do Things with Words*. Harvard University Press, 1962.
- Bach, Kent, and Robert M. Harnish. *Linguistic Communication and Speech Acts*. MIT Press, 1979.
- Bacon, Andrew. *Vagueness and Thought*. Oxford University Press, 2018.
- Barker, Chris. "Clarity and the Grammar of Skepticism." *Mind and Language* 24, no. 3 (2009): 253–73.
- Barker, Chris. "The Dynamics of Vagueness." *Linguistics and Philosophy* 25, no. 1 (2002): 1–36.
- Barker, Chris. "Negotiating Taste." *Inquiry* 56, nos. 2–3 (2013): 240–57.
- Barras, Colin. "We Don't Know Which Species Should Be Classed as 'Human'." BBC News, February 11, 2017. <http://www.bbc.com/earth/story/20160111-what-is-it-that-makes-you-a-human-and-not-something-else>.
- Bierwisch, Manfred. "The Semantics of Gradation." In *Dimensional Adjectives*, edited by Manfred Bierwisch and Ewald Lang. Springer, 1989.
- Braun, David, and Theodore Sider. "Vague, so Untrue." *Noûs* 41, no. 2 (2007): 133–56.
- Brown, Mike. *How I Killed Pluto: And Why It Had It Coming*. Spiegel & Grau, 2010.
- Buchanan, Ray. "A Puzzle about Meaning and Communication." *Noûs* 44, no. 2 (2010): 340–71.
- Burge, Tyler. "Individualism and the Mental." *Midwest Studies in Philosophy* 4, no. 1 (1979): 73–122.

- Burge, Tyler. "Reference and Proper Names." *Journal of Philosophy* 70, no. 14 (1973): 425–39.
- Burnett, Heather. *Gradability in Natural Language: Logical and Grammatical Foundations*. Oxford University Press, 2017.
- Bybee, Jay S. "Memorandum for John Rizzo, Acting General Counsel of the Central Intelligence Agency: Interrogation of Al Qaeda Operative," August 1, 2002. <https://www.justice.gov/sites/default/files/olc/legacy/2010/08/05/memo-bybee2002.pdf>.
- Caie, Michael. "Semantic Indecision." *Philosophical Perspectives* 32, no. 1 (2018): 108–43.
- Cappelen, Herman, and Ernie Lepore. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Blackwell, 2005.
- Cariani, Fabrizio, and Paolo Santorio. "Will Done Better: Selection Semantics, Future Credence, and Indeterminacy." *Mind* 127, no. 505 (2018): 129–65.
- Cartwright, Richard. "Propositions." In *Analytical Philosophy*, edited by R. J. Butler, Vol. 1. Blackwell, 1962.
- Cresswell, Max J. "The Semantics of Degree." In *Montague Grammar*, edited by Barbara Partee. Academic Press, 2014.
- Cumming, Sam. "Variabilism." *Philosophical Review* 117, no. 4 (2008): 525–54.
- Davidson, Donald. "A Nice Derangement of Epitaphs." In *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, edited by Ernest Lepore. Blackwell, 1986.
- Deal, Amy Rose. *A Theory of Indexical Shift: Meaning, Grammar, and Crosslinguistic Variation*. MIT Press, 2020.
- DeRose, Keith. "Gradable Adjectives: A Defence of Pluralism." *Australasian Journal of Philosophy* 86, no. 1 (2008): 141–60.
- DeRose, Keith. "Single-Scoreboard Semantics." *Philosophical Studies* 119 (2004): 1–21.
- Dreier, James. "Relativism (and Expressivism) and the Problem of Disagreement." *Philosophical Perspectives* 23, no. 1 (2009): 79–110.
- Dummett, Michael. *Frege: Philosophy of Language*. Harvard University Press, 1973.
- Dummett, Michael. "Truth." *Proceedings of the Aristotelian Society* 59, no. 1 (1959): 141–62.

- Edmonds, David. "Cake or Biscuit: Why Jaffa Cakes Excite Philosophers." BBC News, February 20, 2017. <http://www.bbc.com/news/magazine-38985820>.
- Egan, Andy, John Hawthorne, and Brian Weatherson. "Epistemic Modals in Context." In *Contextualism in Philosophy*, edited by Gerhard Preyer and Georg Peter. Oxford University Press, 2005.
- Égré, Paul, Benjamin Spector, Adèle Mortier, and Steven Verheyen. "On the Optimality of Vagueness: 'Around,' 'Between' and the Gricean Maxims." *Linguistics and Philosophy* 46, no. 5 (2023): 1075–130.
- Fara, Delia Graff. "Shifting Sands: An Interest-Relative Theory of Vagueness." *Philosophical Topics* 28, no. 1 (2000): 45–81.
- Feferman, Solomon. "Why the Programs for New Axioms Need to Be Questioned." *Bulletin of Symbolic Logic* 6, no. 4 (2000): 401–13.
- Field, Hartry. "Theory Change and the Indeterminacy of Reference." *Journal of Philosophy* 70 (1973): 462–81.
- Fine, Kit. "Vagueness, Truth and Logic." *Synthese* 30, nos. 3–4 (1975): 265–300.
- Fintel, Kai von, and Anthony S. Gillies. "'Might' Made Right." In *Epistemic Modality*, edited by Andy Egan and Brian Weatherson. Oxford University Press, 2011.
- Fodor, Jerry, and Ernie Lepore. "What Cannot Be Evaluated Cannot Be Evaluated and It Cannot Be Supervalued Either." *Journal of Philosophy* 93, no. 10 (1996): 516–35.
- "40 Wall Street." *Wikipedia*, 2020. https://en.wikipedia.org/w/index.php?title=40_Wall_Street&oldid=941342831.
- Frege, Gottlob. *Collected Papers on Mathematics, Logic, and Philosophy*. Edited by Brian McGuinness. Basil Blackwell, 1984.
- Frege, Gottlob. *Philosophical and Mathematical Correspondence*. Edited by Brian McGuinness. Translated by Hans Kaal. University of Chicago Press, 1980.
- Frege, Gottlob. *Posthumous Writings*. Edited by Friedrich Kaulbach, Hans Hermes, and Friedrich Kambartel. University of Chicago Press, 1979.
- Gibbard, Allan. *Thinking How to Live*. Harvard University Press, 2003.
- Gibbard, Allan. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Harvard University Press, 1992.

- Glanzberg, Michael. "Explanation and Partiality in Semantic Theory." In *Metasemantics: New Essays on the Foundations of Meaning*, edited by Alexei Burgess and Brett Sherman. Oxford University Press, 2014.
- Grice, H. P. "Meaning." *Philosophical Review* 66, no. 3 (1957): 377–88.
- Grice, H. P. *Studies in the Way of Words*. Harvard University Press, 1989.
- Heim, Irene, and Angelika Kratzer. *Semantics in Generative Grammar*. Blackwell, 1998.
- Hellan, Lars. *Towards an Integrated Analysis of Comparatives*. Narr, 1981.
- Higginbotham, James. "Truth and Understanding." *Philosophical Studies* 65, no. 1/2 (1992): 3–16.
- Joyce, James M. "A Defense of Imprecise Credences in Inference and Decision Making." *Philosophical Perspectives* 24, no. 1 (2010): 281–323.
- Kaplan, David. "Afterthoughts." In *Themes from Kaplan*, edited by Joseph Almog, John Perry, and Howard Wettstein. Oxford University Press, 1989.
- Kaplan, David. "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals." In *Themes from Kaplan*, edited by Joseph Almog, John Perry, and Howard Wettstein. Oxford University Press, 1989.
- Kaplan, David. "Dthat." In *Syntax and Semantics, Vol. 9: Pragmatics*, edited by Peter Cole. Academic Press, 1978.
- Keefe, Rosanna. *Theories of Vagueness*. Cambridge University Press, 2000.
- Kennedy, Christopher. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30, no. 1 (2007): 1–45.
- Kennedy, Christopher, and Louise McNally. "Scale Structure, Degree Modification, and the Semantics of Gradable Predicates." *Language* 81, no. 2 (2005): 345–81.
- Kennedy, Christopher, and Malte Willer. "Familiarity Inferences, Subjective Attitudes and Counterstance Contingency: Towards a Pragmatic Theory of Subjective Meaning." *Linguistics and Philosophy* 45, no. 6 (2022): 1395–445.
- Khoo, Justin. "Quasi Indexicals." *Philosophy and Phenomenological Research* 100, no. 1 (2020): 26–53.

- Khoo, Justin, and Joshua Knobe. "Moral Disagreement and Moral Semantics." *Noûs* 52, no. 1 (2018): 109–43.
- King, Jeffrey C. *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context*. Oxford University Press, 2021.
- King, Jeffrey C. "The Metasemantics of Contextual Sensitivity." In *Metasemantics*, edited by Alexis Burgess and Brett Sherman. Oxford University Press, 2014.
- King, Jeffrey C. "Strong Contextual Felicity and Felicitous Underspecification." *Philosophy and Phenomenological Research* 97, no. 3 (2018): 631–57.
- Klein, Ewan. "Comparatives." In *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, edited by Arnim von Stechow and Dieter Wunderlich. Walter de Gruyter, 1991.
- Klein, Ewan. "A Semantics for Positive and Comparative Adjectives." *Linguistics and Philosophy* 4 (1980): 1–45.
- Kocurek, Alexander. "Verbal Disagreement and Semantic Plans." *Erkenntnis* 90 (2023): 1061–94.
- Kocurek, Alexander W., Ethan Jerzak, and Rachel Etta Rudolph. "Against Conventional Wisdom." *Philosophers' Imprint* 20, no. 22 (2020): 1–27.
- Lassiter, Daniel, and Noah D. Goodman. "Adjectival Vagueness in a Bayesian Model of Interpretation." *Synthese* 194, no. 10 (2017): 3801–36.
- Lepore, Ernest. "What Model Theoretic Semantics Cannot Do?" *Synthese* 54, no. 2 (1983): 167–87.
- Levi, Isaac. "On Indeterminate Probabilities." *Journal of Philosophy* 71, no. 13 (1974): 391–418.
- Lewis, David. "General Semantics." *Synthese* 22, nos. 1–2 (1970): 18–67.
- Lewis, David. "Index, Context, and Content." In *Philosophy and Grammar*, edited by Stig Kanger and Sven Öhman. Reidel, 1980.
- Lewis, David. "Languages and Language." In *Language, Mind, and Knowledge*, edited by Keith Gunderson. University of Minnesota Press, 1975.
- Lewis, David. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8, no. 3 (1979): 339–59.
- Ludlow, Peter. *Living Words: Meaning, Underdetermination and the Dynamic Lexicon*. Oxford University Press, 2014.

- MacFarlane, John. "Andrew Bacon, *Vagueness and Thought*." *Philosophical Review* 129, no. 1 (2020): 153–58.
- MacFarlane, John. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford University Press, 2014.
- MacFarlane, John. "Epistemic Modals Are Assessment Sensitive." In *Epistemic Modality*, edited by Andy Egan and Brian Weatherson. Oxford University Press, 2011.
- MacFarlane, John. "Future Contingents and Relative Truth." *Philosophical Quarterly* 53, no. 212 (2003): 321–36.
- MacFarlane, John. "Fuzzy Epistemicism." In *Cuts and Clouds: Vagueness, Its Nature, and Its Logic*, edited by Richard Dietz and Sebastiano Moruzzi. Oxford University Press, 2010.
- MacFarlane, John. "Lecture I: Vagueness and Communication." *Journal of Philosophy* 117, no. 11/12 (2020): 593–616.
- MacFarlane, John. "Lecture II: Seeing Through the Clouds." *Journal of Philosophy* 117, no. 11/12 (2020): 617–42.
- MacFarlane, John. "Lecture III: Indeterminacy as Indecision." *Journal of Philosophy* 117, no. 11/12 (2020): 643–67.
- MacFarlane, John. "The Things We (Sorta Kinda) Believe." *Philosophy and Phenomenological Research* 73, no. 1 (2006): 218–24.
- MacFarlane, John. "Vagueness as Indecision." *Aristotelian Society Supplementary Volume* 90, no. 1 (2016): 255–83.
- McGee, Vann, and Brian McLaughlin. "Logical Commitment and Semantic Indeterminacy: A Reply to Williamson." *Linguistics and Philosophy* 27, no. 1 (2004): 123–36.
- McGee, Vann, and Brian McLaughlin. "Review of Timothy Williamson, *Vagueness*." *Linguistics and Philosophy* 21, no. 2 (1998): 221–35.
- McLean, A. N., and P. D. McGreevy. "Ethical Equitation: Capping the Price Horses Pay for Human Glory." *Journal of Veterinary Behavior: Clinical Applications and Research* 5, no. 4 (2010): 203–9.
- Monger, Christopher. *The Englishman Who Went up a Hill but Came Down a Mountain: A Novel*. Miramax Books, 1995.
- Ninan, Dilip. "Counterfactual Attitudes and Multi-Centered Worlds." *Semantics and Pragmatics* 5 (2012): 1–57.
- Ninan, Dilip. "Semantics and the Objects of Assertion." *Linguistics and Philosophy* 33, no. 5 (2010): 335–80.

- Nowak, Ethan. "No Context, No Content, No Problem." *Mind & Language* 36, no. 2 (2021): 189–220.
- Nunberg, Geoffrey. "The Social Life of Slurs." In *New Work on Speech Acts*. Oxford University Press, 2017.
- Perry, John. "Frege on Demonstratives." *Philosophical Review* 86, no. 4 (1977): 474–97.
- Plunkett, David, and Timothy Sundell. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosopher's Imprint* 13 (2013): 1–37.
- Plunkett, David, and Timothy Sundell. "Metalinguistic Negotiation and Matters of Language: A Response to Cappelen." *Inquiry* 68, no. 10 (2021): 3123–47.
- Plunkett, David, and Timothy Sundell. "Metalinguistic Negotiation and Speaker Error." *Inquiry* 64 (2021): 142–67.
- Price, Huw. "Does 'Probably' Modify Sense?" *Australasian Journal of Philosophy* 61, no. 4 (1983): 396–408.
- Putnam, Hilary. "The Meaning of 'Meaning'." *Minnesota Studies in the Philosophy of Science* 7 (1975): 131–93.
- Putnam, Hilary. *Representation and Reality*. MIT Press, 1988.
- Qing, Cyang, and Michael Franke. "Gradable Adjectives, Vagueness, and Optimal Language Use: A Speaker-Oriented Model." *Proceedings of SALT* 24 (2014): 23–41.
- Rabern, Brian. "Against the Identification of Assertoric Content with Compositional Value." *Synthese* 189, no. 1 (2012): 75–96.
- Rabern, Brian. "Semantic Monsters." In *Routledge Handbook of Linguistic Reference*, edited by Heimir Geirsson and Stephen Biggs. Routledge, 2021.
- Raffman, Diana. *Unruly Words: A Study of Vague Language*. Oxford University Press, 2014.
- Raffman, Diana. "Vagueness Without Paradox." *Philosophical Review* 103, no. 1 (1994): 41–74.
- Recanati, François. *Literal Meaning*. Cambridge University Press, 2003.
- Richard, Mark. *When Truth Gives Out*. Oxford University Press, 2008.
- Rinard, Susanna. "A Decision Theory for Imprecise Probabilities." *Philosophers' Imprint* 15, no. 7 (2015): 1–16.

- Roberts, Craige. "Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics." *Semantics and Pragmatics* 5, no. 6 (2012): 1–69.
- Russell, Bertrand. "Vagueness." *Australasian Journal of Philosophy and Psychology* 1, no. 2 (1923): 84–92.
- Santorio, Paolo. "Context-Free Semantics." *Oxford Studies in Philosophy of Language* 1 (2019): 208–39.
- Santorio, Paolo. "Reference and Monstrosity." *Philosophical Review* 121, no. 3 (2012): 359–406.
- Schiffer, Stephen. *The Things We Mean*. Oxford University Press, 2003.
- Schiffer, Stephen. "Two Issues of Vagueness." *The Monist* 81, no. 2 (1998): 193–214.
- Schroeder, Mark. *Being for: Evaluating the Semantic Program of Expressivism*. Oxford University Press, 2008.
- Scontras, Gregory, Michael Henry Tessler, and Michael Franke. "A Practical Introduction to the Rational Speech Act Modeling Framework," 2021. <https://doi.org/https://doi.org/10.48550/arXiv.2105.09867>.
- Searle, John. "Literal Meaning." *Erkenntnis* 13, no. 1 (1978): 207–24.
- Seuren, Pieter A. M. "The Comparative." In *Generative Grammar in Europe*, edited by Ferenc Kiefer and Nicolas Ruwet. Reidel, 1973.
- Silk, Alex. *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. Oxford University Press, 2016.
- Soames, Scott. *Understanding Truth*. Oxford University Press, 1999.
- Soria-Ruiz, Andrés. "Is Metalinguistic Usage a Conversational Implicature?" *Topoi* 42 (2023): 1027–38.
- Stalnaker, Robert. "Assertion." In *Syntax and Semantics, Vol. 9: Pragmatics*, edited by P. Cole. Academic Press, 1978.
- Stalnaker, Robert. "Common Ground." *Linguistics and Philosophy* 25, nos. 5–6 (2002): 701–21.
- Stalnaker, Robert. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press, 1999.
- Stanley, Jason. "Modality and What Is Said." *Language and Mind* 16 (2002): 321–44.
- Stechow, Arnim von. "Comparing Semantic Theories of Comparison." *Journal of Semantics* 3 (1984): 1–77.

- Suikkanen, Jussi. "Contextualism, Moral Disagreement, and Proposition Clouds." Edited by Russ Shafer-Landau. *Oxford Studies in Metaethics* 14 (2019): 47–69.
- Sullivan, Ed. "Movie Review: Ffynnon Garw." *Professional Surveyor* 18, no. 8 (1998). <https://archives.profsurv.com/magazine/article.aspx?i=368>.
- Sundell, Timothy. "The Tasty, the Bold, and the Beautiful." *Inquiry* 59, no. 6 (2016): 793–818.
- Tappenden, Jamie. "Some Remarks on Vagueness and a Dynamic Conception of Language." *Southern Journal of Philosophy* 33, no. supplement (1994): 193–201.
- The Englishman Who Went up a Hill but Came Down a Mountain*. Miramax Films, 1995.
- Thomason, Richmond. "Introduction." In *Formal Philosophy: Selected Papers of Richard Montague*, edited by Richmond Thomason. Yale University Press, 1974.
- Travis, Charles. "Pragmatics." In *A Companion to the Philosophy of Language*, edited by Bob Hale and Crispin Wright. Blackwell, 1997.
- Voltaire. *Dictionnaire philosophique*. Vol. 4. Oeuvres Complètes. Hachette, 1960.
- Williamson, Timothy. "Soames on Vagueness." *Philosophy and Phenomenological Research* 65, no. 2 (2002): 422–28.
- Williamson, Timothy. *The Philosophy of Philosophy*. Blackwell, 2007.
- Williamson, Timothy. *Vagueness*. Routledge, 1994.
- Wilson, Mark. "Predicate Meets Property." *Philosophical Review* 91, no. 3 (1982): 549–89.
- Wittgenstein, Ludwig. *Philosophical Investigations*. 2nd ed. Blackwell, 1958.
- Wright, Crispin. "Language Mastery and the Sorites Paradox." In *Truth and Meaning*, edited by Gareth Evans and John McDowell. Oxford University Press, 1976.
- Yalcin, Seth. "Dynamic Semantics." In *Routledge Companion to Philosophy of Language*, edited by Gillian Russell and Delia Graff Fara. Routledge, 2012.
- Yalcin, Seth. "Epistemic Modals." *Mind* 116 (2007): 983–1026.

Index

- α (parameter in the Rational Speech Act model), 176
- | (domain restriction), 104
- [[]] (denotation), 103
- accommodation, 39, 41
- agreement, 144
- ambivalence, 63
- anti-individualism, 11–12, 146–147, 154
 - and demonstratives, 14
 - and vague language, 15
- approximately true, 38
- argument from disagreement *See* disagreement
- Assert, 46
- Assertoric Compatibility, 141
- assertoric content, 98
- assignment, 117
 - constraints on, 119
 - of the context, 119
- Bacon, Andrew, 20
- Barker, Chris, 55, 61–62, 66
- Bedeutung*, 1
- betting behavior
 - and vagueness-related ambivalence, 83
- bifurcated common ground *See* common ground
- bivalence, 67–72
- Bivalence 1, 67
- Bivalence 2, 68
- blind externalism, 153–154
- borderline
 - taking something to be, 80
- Braun, David, 32, 36–43
- Buchanan, Ray, 32–33, 37, 43–45
- Burge, Tyler, 154
- Classical Contents, 31
- Classical Pragmatics, 31
- cloudy picture, 31
- common ground, 22, 76
 - bifurcated, 39, 41, 161
- common knowledge
 - introspection principles for, 77–78
- communication, 5
- comparatives
 - semantics of, 103
- compared to α , 104
- comparison class *See* reference class
- compositional semantic value, 98
- compositional semantics, 99
- concepts, 79

- Confirm/Deny, 48
 consequence, 100
 constraints, 135
 content
 and intension, 155
 as constraint, 157
 in negotiated contextualism, 165
 individuation of, 121–123
 information, 173
 context, 98
 abstract vs. concrete, 108
 cloud of, 34
 indeterminacy of, 108
 context change potential, 121
 context set, 23
 contextualism, 67, 164 *See also*
 single-scoreboard semantics,
 discourse contextualism,
 negotiated contextualism
 convention shifting, 129–131
 conversational score, 164
 coordination, 13, 79
 credence, 62
 in vague propositions, 84

 Dandele, Sophia, 71
 Davidson, Donald, 135–136
 degree function, 103
 degree semantics
 for gradable adjectives, 101–104,
 183
 delineation, 54
 demonstratives, 12–14
 DeRose, Keith, 160–162
 diagonal proposition, 23
 diagonalization, 22–27, 55, 131–133
 disagreement, 140–142
 argument from, 138, 142, 160
 locus of, 141
 merely verbal, 86, 142
 metalinguistic, 142
 primary, 171
 discourse contextualism, 162–164
 distribution data, 102
 dynamic semantics, 123–125

 Égré, Paul, 187–192
 entailment, 100
 epistemic modals, 34, 45–51
 epistemic shift, 125–128
 epistemicism, 8–12, 27, 85, 153
 and partial belief, 62–63
 error theory, 148–149
 ESMV model, 187
 excluded middle, 68–72
 Expressive Incompatibility, 141
 expressivism
 negation problem for, 70
 plan, 64–67
 externalism 154 *See also* blind
 externalism

 Fara, Delia Graff, 18
 felicitous underspecification, 21–22
 Fintel, Kai von, 33–36, 45–54
 for an *F*, 104
 Frege's puzzle, 127
 Frege, Gottlob, 1–2, 6
 fuzzy semantics, 20, 81

 The Garth, 74
 Gibbard, Allan, 64–65, 70, 71, 152
 Gillies, Anthony, 33–36, 45–54
 Glanzberg, Michael, 114–115
 Goodman, Noah D., 182–187
 gradable adjectives, 101
 relative vs. absolute, 104
 scale structure of, 104
 Grice, H. P., 44

 Heim, Irene, 123–124
 higher-order vagueness, 76–78
 homonyms, 159

- homophonic clauses
 - in semantics, 114
- hypercontexts, 106, 117
 - and assertoric content, 107
- hyperconventions, 129
- hyperplans, 64
 - and indifference, 68–72
 - in the index, 106
- I (first-person pronoun), 125–127
- $I(u, s)$ (informativeness of an utterance), 176
- idempotence, 189
- indecision, 63
 - vs. indifference, 70
 - vs. uncertainty, 82
- index, 98
- indifference
 - vs. indecision, 70
- inexact knowledge, 62
- information
 - probabilistic, 173
- information contents, 173
- informational model of assertion, 173
- informativeness
 - of an utterance, 176, 185, 187
 - of vague utterances, 182
- interest-relative, 67
- introspection *See* common
 - knowledge, Positive
 - Introspection, Negative
 - Introspection
- Jerzak, Ethan, 129–131
- Kaplan, David, 126–127
- Kennedy, Chris, 101–104
- Khoo, Justin, 164–170
- Khoo-content, 165
- King, Jeffrey, 21–22, 29, 87–93
- KJR, 129
- Kocurek, Alexander W., 129–131
- Kullback-Leibler divergence, 178, 187, 191
- Lassiter, Daniel, 182–187
- law of excluded middle *See*
 - excluded middle
- Lepore, Ernie, 134
- lexical underspecification, 111–113, 118
- metalinguistic negotiation, 141
- mixed cases, 81, 91, 150–152, 157–158, 161, 163
- mode of presentation, 8
- monsters, 126
- Montague semantics
 - vs. Davidsonian semantics, 134–135
- Montague, Richard, 134
- Mortier, Adèle, 187–192
- negation problem, 70
- Negative Introspection, 77
- negotiated contextualism, 164–170
- nonfactual parameter, 39, 55
- panvariabilism, 116–121, 135–136
- partial belief 62, 80 *See also*
 - vagueness-related partial belief
 - standard partial belief
- planning states, 64
- Plunkett, David, 137, 138
- POS*, 102
 - semantics of, 103
- Positive Introspection, 77
- $Poss_i$, 34
- postsemantics, 99, 119
- Potts, Chris, 33
- pragmatic listener, 183
- principle of excluded middle *See*
 - excluded middle

- prior and passing theories, 135
- priors
 - over thresholds, 193
- probabilities
 - imprecise, 84
 - of vague propositions, 82
- probably, 50
- proposition, 36
 - cloud of, 30
- proposition type, 43
- Putnam, Hilary, 154
- quasi-indexical, 165, 170
- question under discussion, 87
- question under discussion (QUD), 178
- reference class, 15–18, 104
- RSA model, 175–180
- Rudolph, Rachel Etta, 129–131
- ruling out, 157
- s_c (threshold function), 103
- Schiffer, Stephen, 80–82
- semantic minimalism, 15
- semantic shift, 128–131
- semantic value *See* compositional semantic value
- semantics, 96–101
- shared meaning task, 153
- Sider, Theodore, 32, 36–43
- Silk, Alex, 162–164
- Simple Model of Communication, 5, 31
- single-scoreboard semantics, 160–162
- soft-max, 176
- sorites paradox, 72–75
- SPB *See* standard partial belief
- speaker-meaning, 145
- Spector, Benjamin, 187–192
- Stalnaker, Robert, 22–23, 44, 59, 146
- standard partial belief, 81–82
- Strong Kleene logic, 71
- Sundell, Tim, 138
- Sundell, Timothy, 137
- superlatives
 - semantics of, 103
- supervaluationism, 42
 - and the sorites, 72
- Thomason, Richmond, 134
- three-valued logic, 71
- threshold, 103
- tolerance, 20, 73
- transparency of meaning
 - intentions, 44
- transparent (truth predicate), 167
- Travis, Charles, 112–113
- truth
 - at a context, 67, 99
 - at a context and index, 99
 - of an utterance, 68
- truth predicate
 - vagueness of, 68
- truth-conditional semantics, 98
- truth-value gaps, 67
- type-token ambiguity, 88
- uncertainty, 63
- understandings, 112
- Uniformity, 23
- update value, 165
- uptake, 59
- vague thought, 78–85
- vagueness
 - utility of, 85–86, 181
- vagueness-related partial belief, 81–82
- variabilism, 118

- variables, 117
 - in mathematics, 159
- Verheyen, Steven, 187–192
- Voltaire, 86
- VPB *See* vagueness-related partial belief
- Williamson, Timothy, 68
- Wittgenstein, Ludwig, 153
- world/assignment pairs, 121
 - in dynamic semantics, 123
- world/hyperplan pairs, 64