

# Living Words



# Living Words

*Meaning Underdetermination  
and the Dynamic Lexicon*

Peter Ludlow

**OXFORD**  
UNIVERSITY PRESS

**OXFORD**

UNIVERSITY PRESS

Great Clarendon Street, Oxford, ox2 6DP,  
United Kingdom

Oxford University Press is a department of the University of Oxford.  
It furthers the University's objective of excellence in research, scholarship,  
and education by publishing worldwide. Oxford is a registered trade mark of  
Oxford University Press in the UK and in certain other countries

© Peter Ludlow 2014

The moral rights of the author have been asserted

First Edition published in 2014

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in  
a retrieval system, or transmitted, in any form or by any means, without the  
prior permission in writing of Oxford University Press, or as expressly permitted  
by law, by licence or under terms agreed with the appropriate reprographics  
rights organization. Enquiries concerning reproduction outside the scope of the  
above should be sent to the Rights Department, Oxford University Press, at the  
address above

You must not circulate this work in any other form  
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press  
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number:

ISBN 978-0-19-871205-3

Printed and bound in Great Britain by  
CPI Group (UK) Ltd, Croydon, CR0 4YY

Links to third party websites are provided by Oxford in good faith and  
for information only. Oxford disclaims any responsibility for the materials  
contained in any third party website referenced in this work.

*For my mother, who litigated word meanings  
with me all the time*

Because, Soferim Bebel, if it goes to that . . . every person, place and thing in the chaosmos of Alle anyway connected with the gobbly-dumbed turkey was moving and changing every part of the time: the traveling inkhorn (possibly pot), the hare and the turtle pen and paper, the continually more or less intermisunderstanding minds of the anticollaborators, the as time went on as it will variously inflected, differently pronounced, otherwise spelled, changeably meaning vocable scriptsigns. No, so help me Petault, it is not a mis-effectual whyacinthinous riot of blots and blurs and bars and balls and hoops and wriggles and juxtaposed jottings linked by spurts of speed: it only looks as like it as damn it.

(James Joyce, *Finnegan's Wake*)

## Preface

The material in this book presented a challenge for me as an author. On the one hand, the basic ideas of this book—meaning underdetermination, dynamic word meanings, word meaning litigation, and lexical warfare—can be accessible to a general audience when properly explained. On the other hand, these basic ideas present puzzles and worries that quickly lead us into some of the more difficult terrain in contemporary analytic philosophy.

One thought I had was to write two books—one for a general audience and one for a specialized audience trained in logic and the semantics of natural language, but I decided this would not be the best path. In the first place it underestimates the abilities of a non-philosophically trained audience. A good author should be able animate the technical issues and walk such an audience through the puzzles—or at least give it a sense of what the big puzzles are.

In the second place, I believe that it serves technical philosophy well to think about how it fits within a broader conversation and to see that, yes, this technical work does indeed have consequences that nonspecialists can understand and with which they can engage at a high level.

Ultimately, I opted for a single monograph. This required some editorial decisions on my part that bear note. As a general rule, the technical material comes later in the book. I've made an effort to make the technical material accessible as far as possible, but in some cases background in logic and the philosophy of language are necessary for the material to be fully accessible. It is my hope that, even if this material is not completely accessible on the first pass, the reader will at least feel invited to engage the relevant background material and return to these topics at a later time.



## Acknowledgements

Crafting a book aimed at both a technical audience and a general audience tested the outer limits of my abilities as a writer, and insofar as I even had the courage to make the attempt I owe thanks to a number of people who have taught me how to write for a general audience—in particular Peter Catapano of the *New York Times*, Katrina van den Heuvel at *The Nation*, and of course Mark Wallace (aka Walker Spaight), with whom I co-authored *The Second Life Herald: The Virtual Tabloid that Witnessed the Dawn of the Metaverse*. (Peter Catapano also gets credit for helping me come up with the title for this book.)

Turning to matters of content, this book has been in the works for about a decade and I have many people to thank for valuable suggestions and difficult yet very helpful questions over the years. In particular, I would like to thank Josh Armstrong, David Braun, Susan Brennan, Liz Camp, Herman Cappelen, Chris Gauker, Patrick Grim, Gil Harman, Liz Harman, John Hawthorne, Richard Larson, Ernie Lepore, Rebecca Mason, Brian McLaughlin, Francois Recanati, Dan Sperber, Jason Stanley, Matthew Stone, Tim Sundell, Paul Teller, Deirdre Wilson, and David Zarefsky for these helpful discussions.

Additional help came when more or less complete versions of this material were presented in minicourses at Beihan University, Beijing China, August, 2011, and the International Summer School in Semantics and Cognitive Science, Pumpula, Latvia, July 2012.

In addition, smaller portions of this work have been presented in various talks over the past decade. Among those places: the Conference on Cognitive Systems as Representational Systems, Nicolaus Copernicus University, Torun, Poland, 2004; Meaning and Communication Conference, Lisbon, 2005; Mental Matters: The Philosophy of Linguistics, Dubrovnik, 2005; University of Toronto, Dept. of Philosophy, 2005; University of Central Oklahoma, 2006; Context and Communication Conference, University of Oslo, Oslo,

## X ACKNOWLEDGEMENTS

2006; International Conference on Linguistics and Epistemology, University of Aberdeen, Scotland, 2007; American Philosophical Association Central Division Meeting, Chicago, 2008; World Congress of Philosophy, Seoul, 2008; American Philosophical Association Pacific Division Meeting, Vancouver, 2009; Conference on Contextualism and Truth, Arche, University of St Andrews, Scotland, 2009; University of Buenos Aires, 2009; Dept. of Philosophy, UNLV, September 2009; Conference on Contextualism and Compositionality, University of Paris, 2010; Workshop in Semantics and Philosophy of Language, University of Chicago, 2010; Kansas State University, Dept. of Philosophy, 2011; Rutgers AEF Interdisciplinary Meeting on Approaches to Reference, Rutgers University, 2011; International Conference on Language and Value, Beijing Normal University, 2011. I am very grateful to the audiences at those conferences who pushed this work and helped me to develop it in profitable ways.

Portions of this work have previously been published in article form. Parts of section 2.1 appeared in “The Myth of Human Language,” *Croatian Journal of Philosophy*, 2006. Portions of section 5.1 appeared in “Cheap Contextualism,” *Noûs. Philosophical Issues 16: Annual Supplement*, 2008, and bits of 5.3 appeared in “Understanding Temporal Indexicals,” (with reply by John Perry) in M. O’Rourke and C. Washington (eds), *Situating Semantics: Essays on the Philosophy of John Perry* (Cambridge, MA: MIT Press, 2007).

Finally, thanks are due to a pair of OUP anonymous reviewers and to Peter Momtchiloff for editorial guidance and helping me to keep my eye on the ball until this project was completed.

# Contents

1. Introduction	1
1.1 The Static Lexicon vs. the Dynamic Lexicon	1
1.2 Lexical Warfare	7
1.3 Unreflective Entrainment	25
2. Norms of Word Meaning Litigation	39
2.1 Recognizing and Engaging Meaning Mismatch	39
2.2 ‘Planet’	41
2.3 ‘Rape’	51
2.4 ‘Person’	55
2.5 Scalia’s <i>Original Meaning</i> Thesis	64
2.6 True When it was Uttered?	66
3. The Nature of the Dynamic Lexicon	72
3.1 Some Features of the Dynamic Lexicon	72
3.2 Core Ideas	80
3.3 Terminology	84
3.4 Precursors	90
4. Meaning Underdetermination, Logic, and Vagueness	101
4.1 Meaning Underdetermination and Semantics	102
4.2 Word Meaning Modulation and Logic	114
4.3 Vagueness	122
5. Consequences for Analytic Philosophy	129
5.1 Contextualism in Epistemology	129
5.2 Paderewski, Peter, and Pierre	146
5.3 Understanding Indexicals	151
6. Metaphor and Beyond	158
<i>Bibliography</i>	165
<i>Index</i>	177



# 1

## Introduction

### 1.1 The Static Lexicon vs. the Dynamic Lexicon

Quite often people ask me how many books I've written. When they do (for example, on airplanes), I pause and say, "well... it depends on what you mean by 'book.'" I have edited several volumes of previously published work by others. Do these edited volumes count as books? Some people (most non-academics) say yes, and others say no. I have written a couple of eBooks; do they count as books? But wait, one isn't published yet. And the one that is published is only about 50 pages long. Book? Again the answer I get varies. Was my Columbia University dissertation a book? By the way, it was "published," with minor revisions, by the University of Indiana Linguistics Club. Book? The same book? What about drafts of books that are sitting on my hard drive? Are they books? Is a co-authored book a "book I wrote?" It takes a few minutes of asking these questions before I can answer and tell my conversational partner whether I have written two or three or six or ten books.

This story is odd in a way, because 'book' is one of the first words we English speakers learn, and it has been with us for a long time. It comes from the old English 'boc', which seemed to apply to any written document. The shared meaning has evolved over the past thousand years to be somewhat narrower than that (not every written document is a book) and in some ways broader (think eBook) but even after a millennium of shared usage the meaning is quite open-ended. And there are elements of the meaning that can change radically on a conversation-by-conversation basis.

## 2 INTRODUCTION

Far from being the exception, I think this is typical of how things are with the words we use. Even for well-entrenched words their meanings are open ended and can change on the fly as we engage different conversational partners. Consider a word like ‘sport’. Does it include bowling? Mountain Climbing? Darts? Chess? Or consider words like ‘freedom’, ‘journalist’, or (less loftily) ‘sandwich’ and ‘doll’. All of these words have meanings that are underdetermined, and we adjust or modulate their meanings on a conversation-by-conversation basis. Their meanings are *dynamic*.

These facts seem to fly in the face of the traditional view of language, which is more or less the following: Languages like Urdu, German, Polish, and Portuguese are fairly stable abstract systems of communication that are learned (with varying degrees of success) by human beings. Those humans in turn use the languages that they have learned to communicate ideas, perform certain tasks (by giving orders, instructions, etc.), and in some cases as media for artistic expression. It is often supposed that the better one learns a language the better equipped one is to successfully communicate, accomplish complex tasks, etc. Sometimes the standard view uses the metaphor of language as a widely shared common currency that agents use to communicate, with individual words being the common coins of the realm. These common coins are also supposed to be more or less fixed. Of course everyone believes that language undergoes change, but according to the standard view the pace of change is glacial; there is a long slow gradual evolution from Old English to Middle English and on to Contemporary English. On the standard view word meanings change slowly, and the change is largely uniform across the population of language users.

In this book I follow recent work in philosophy, linguistics, and psychology that rejects the standard, static picture of language, and instead highlights the extreme context sensitivity of language. From this alternative point of departure I will develop an alternative *dynamic* theory of the nature of language and the lexicon. This alternative theory will reject the idea that languages are stable abstract objects

that we learn and then use; instead, human languages are things that we build on a conversation-by-conversation basis. We can call these one-off fleeting things *microlanguages*. I will also reject the idea that words are relatively stable things with fixed meanings that we come to learn. Rather, word meanings themselves are dynamic and massively underdetermined.

What do I mean when I say that word meanings are dynamic and underdetermined? First, when I say that the meaning of a term is *dynamic* I mean that the meaning of the term can shift between conversations and even within a conversation. As I noted, everyone agrees that word meanings can shift over time, but I will argue that they also shift as we move from context to context during the day.

These shifts of meaning do not just occur between conversations; I think that they also occur *within* conversations—in fact I believe that conversations are often designed to help this shifting take place. That is, when we engage in conversation, much of what we say does not involve making claims about the world but it involves instructing our communicative partners about how to adjust word meanings for the purposes of our conversation.

For example, the linguist Chris Barker (2002) has observed that many of the utterances we make play the role of shifting the meaning of a predicate. Sometimes when I say “Jones is bald,” I am not trying to tell you something about Jones; I am trying to tell you something about the meaning of ‘bald’—I am in effect saying that for the purposes of our current conversation, the meaning of ‘bald’ will be such that Jones is a safe case of a bald person (more precisely, that he safely falls in the *range* of the predicate ‘bald’) and that from this point forward in the conversation everyone balder than Jones is safely in the range of ‘bald’.<sup>1</sup> Barker’s observation generalizes to a broad class of our linguistic practices; even if it appears that we are making assertions of fact, we are often doing something else altogether. Our utterances

<sup>1</sup> I’ll explain what I mean by ‘range’ in s. 3.2, but a warning to analytic philosophers: It is not quite the same thing as the extension of the predicate; I take an extension to have a fixed membership but a range to be open and underdetermined.

## 4 INTRODUCTION

are *metalinguistic*—we are using our conversation to make adjustments to the language itself, perhaps to clarify the claims that will only follow later.

We have other strategies for shifting word meanings in a conversation. Sometimes we say things like “Well if Jones is bald then Smith is bald.” I think that what is happening when we do this is that we are trying to persuade our interlocutor that, given our agreement that Jones is safely in the range of ‘bald’, Smith ought to be considered safely in the range of ‘bald’ too, or perhaps we are running a *reductio* argument to persuade our interlocutor that Jones shouldn’t count as in the range of ‘bald’.

Why does the difference between this dynamic theory and the standard (relatively static) theory matter? First, while the static theory is not *universally* held (as we will see, a number of contemporary philosophers and linguists have rejected it) it is at least *widely* held by both academics and non-academics, ranging from philosophers and language instructors, to anthropologists and computational linguists, to politicians and political pundits. Second, even though the standard theory is not universally accepted, the basic assumptions of the standard view have nevertheless crept into the way problems are tackled in all of these domains—sometimes with devastating consequences.

For example, the standard view has led anthropologists and psychologists to think that languages constrain the conceptual space of language users. It has led to wooden approaches to language instruction on the one hand and to failed attempts at human/machine communication on the other. On the political end, it has led to silliness on both the left and the right by way of attempts to clean up or reform or otherwise render standard languages politically correct—a general sentiment that has led to downright discriminatory social policies like English Only laws and, in its extreme form, to attempts at language purification by Fascists like Mussolini.

Finally, I believe that the standard view has led to imbroglios in contemporary analytic philosophy on topics ranging from the theory of sense and reference, to the philosophy of time, skepticism in epistemology, and the problem of vagueness. To see our way out of these

imbroglios we need to attend to the more accurate picture of the nature of language as a dynamic object. That is, it is not enough to pay lip service to the idea that language is dynamic; we have to ensure that static assumptions have not crept into our philosophical theorizing. Static assumptions need to be isolated and removed if we want to avoid philosophical conundrums.

For example, as I will argue in section 5.1, the meaning of the term ‘know’ can shift from conversational context to conversational context. Someone might ask me if I know where the car keys are, and I may truly say yes, even though in an epistemology class I might say that I can’t be sure that car keys and cars even exist (I could be a brain in a vat, after all). How can I know where my keys are if I don’t even know they exist? One way of understanding what is going on here is to say that the meaning of ‘know’ has shifted between its use in the epistemology class and its use in an everyday context. The meaning of ‘knowledge’ in an epistemology class is much more stringent than the meaning of ‘knowledge’ in everyday contexts. There are countless examples of this sort of phenomenon. Every field has terms that get specialized meanings when people are talking shop. For example, the materials scientist will say that the glass in a window pane is liquid when she is wearing her scientist hat, but presumably will not call it a liquid in everyday conversation.

Word meanings are dynamic, but they are also *underdetermined*.<sup>2</sup> What this means is that there is no complete answer to what does and doesn’t fall within the range of a predicate like ‘red’ or ‘bald’ or ‘hexagonal’ (yes, even ‘hexagonal’). We may sharpen the meaning and we may get clearer on what falls in the range of these predicates (and we may willingly add or subtract individuals from the range), but we never completely sharpen the meaning and we never completely nail down the extension of a predicate. For example, we might agree that Jones

<sup>2</sup> I believe this notion is similar to Friedrich Waismann’s idea of meanings being “open textured,” as developed in Shapiro (2006) and also Gauker (2013). Both of these works came to my attention after I completed the bulk of this work and I haven’t had an opportunity to study these proposals in detail.

## 6 INTRODUCTION

is safely in the range of ‘bald’, but there are still many cases where the meaning of ‘bald’ isn’t fixed. We haven’t fixed the meaning of ‘bald’ for people with more hair than Jones, or for people with about the same amount of hair as Jones but distributed differently, or for people who shave their heads, or for nonhumans, etc.

Some theorists think that there is a core meaning for a term that is the absolute sense of the term but that we are pragmatically licensed to use the term loosely. So, for example, ‘bald’ means absolutely bald—not one single hair,<sup>3</sup> ‘flat’ means absolutely flat, etc. There are various ways of executing this idea. For example Laserson (1990) has talked of “pragmatic halos” surrounding the core, absolute sense of the terms; Recanati (2004) and Wilson and Carston (2007) have argued that we begin with the absolute meaning and are “pragmatically coerced” to modulate to less precise meanings. I don’t believe this view is correct. In this book I will argue that the “absolute” sense of a term (if it even exists) is not privileged but is simply one modulation among many—there is no core or privileged modulation.

This isn’t just the case for predicates like ‘bald’ but, I will argue, all predicates, ranging from predicates for things like ‘person’ and ‘tree’, predicates for abstract ideas like ‘art’ and ‘freedom’, and predicates for crimes like ‘rape’ and ‘murder’. You may think that there is a core, fully fleshed out meaning that these predicates refer to, but you would be quite mistaken—even in the legal realm the meanings are not fully fleshed out, not by *Black’s Law Dictionary*, nor by written laws, nor by the intentions of the lawmakers and founding fathers.<sup>4</sup> Indeed, I would argue that this is also the case with mathematical and logical predicates like ‘straight line’ and ‘entailment’. The meanings of all these predicates remain open to some degree or other, and are sharpened as needed when we make advances in mathematics and logic.

You might think that underdetermined meanings are defective or inferior and perhaps things to be avoided, but in my view they can’t

<sup>3</sup> Of course on this view one presumably needs some absolute sense of ‘hair’, which I think would be difficult to spell out. Is one cell of hair DNA in a hair follicle a hair?

<sup>4</sup> See Endicott (2000) for discussion.

be avoided (even in mathematical and logical cases), and in any case there is no point in avoiding them since we reason perfectly well with words having underdetermined meanings. I will attempt to show how this works and in particular how we can have a formal semantics of natural language even though we are admitting massive meaning underdetermination. The received wisdom seems to be that semantics demands precision and fully determinate meanings. Whatever the merits of precision and fully determinate meanings, semantics has no need for them.

Finally, we will see that the static view has infected analytic philosophy, with the result that philosophy has accumulated a number of seemingly intractable puzzles that, I believe, all have their roots in these two errors—the assumption that the lexicon is static and that meanings are fully determined. I'll give a handful of examples of where this has taken place, but it is my belief that once we pull on these threads many more puzzles in contemporary philosophy will begin to unravel.

## 1.2 Lexical Warfare

As we will see, in certain cases meaning modulation is automatic, and to some degree cooperative. But there are also cases in which we are aware that meaning modulation is taking place—not only aware, but actually engaged in finding ways to litigate for our preferred modulation.

'Lexical warfare' is a phrase that I like to use for battles over how a term is to be understood. Our political discourse is full of such battles; it is pretty routine to find discussions of who gets to be called 'Republican' in the United States (as opposed to RINO—Republican in Name Only), what 'freedom' should mean, what gets called 'rape', and the list goes on.

Lexical warfare is important because it can be a device to marginalize individuals within their self-identified political affiliation (e.g. making them not true Republicans), or it can beguile us into ignoring true threats to freedom (e.g. by focusing on threats from government while being blind to threats from corporations, religion, and custom),

## 8 INTRODUCTION

and in cases like ‘rape’ the definition can have far-reaching consequences for social policy (we will discuss this case in Chapter 2).

Lexical warfare is not exclusively concerned with how terms are to be defined—it can also work to attach either a negative or positive aspect to a term. So, famously, Ronald Reagan successfully attached a negative patina to ‘liberal’, while a term like ‘patriot’ has a positive affect (few today reject the label ‘patriotic’, they rather argue for why they are entitled to it).

A good example of the concern for affect in lexical warfare can be found in an amicus brief written on behalf of Andrew Auernheimer, who is better known under his hacker *nom de guerre*, ‘weev’. In 2013 weev was sentenced to forty-one months in jail for (with a friend) using a script to harvest information that AT&T had left on unprotected web pages. The amicus brief, filed by the Mozilla Foundation and a number of computer scientists, security, and privacy experts, raised a number of issues why weev’s actions should not be considered illegal (and indeed, argued that they were routine actions for security professionals). It also raised an issue about the commonly used phrase ‘brute force method’—a common expression in computer science for methods that exhaustively evaluates all possible solutions (for example, a brute force method in a chess program would work through the outcome of every possible combination of moves rather than construct a heuristic strategy). As the amici observed in a section titled “1. ‘Brute force’ is not nefarious,” the affect normally attaching to “brute force’ should be detached in this context.<sup>5</sup>

The government may refer to the “account slurper” as a “brute force” technique. That term has a particular and innocuous meaning: an approach to a problem that “evaluat[es] all possible solutions.” Alfred V. Aho, *Complexity Theory*, in *Computer Science: The Hardware, Software and Heart of It* 241, 257 (Edward K. Blum & Alfred V. Aho eds., 2011). Despite the thuggish name, there is nothing nefarious about using a “brute force” technique to solve a problem.

<sup>5</sup> <<http://torekeland.com/wp-content/uploads/2013/07/Mozilla-Amicus.pdf>> (last accessed July 2013).

We will get back to the role of meaning modulation in controlling affect in a bit, but first I want to point out a range of cases of meaning negotiation, just so we have some idea of the scope of the phenomenon.

*‘Doll’*

When I was in third grade (in 1965) I received a toy that was designed to look like a native American, and which was called “Chief Cherokee”. The nonstandard thing about the toy was its size, which was about that of a Barbie or Ken doll. One day my father came home from work and asked “why are you playing with a doll” to which I objected that it was not a doll. The term “action figure” had not been invented yet (or at least I hadn’t heard it) so I was left just calling it a toy—no doll. An argument ensued, but I don’t remember the particulars.

Apart from what this story tells us about the socialization of boys in 1960s America, it actually points to a really interesting question—just what kinds of things are in the range of ‘doll’? It clearly has nothing to do with the material substrate; there is a long tradition of making dolls from cornhusks and socks and presumably anything, and gender doesn’t seem to matter (cf. Ken dolls) and dolls do seem to come in every possible size (consider how small Russian “babushka” dolls can get). My view, of course, is that the definition is open-ended and dynamic and we can play with it as suits our purposes. Or we can try to if our interlocutors are willing to go along with us.

For the most part how we define ‘doll’ doesn’t have important consequences. But there are exceptions. For example, until recently doll imports were taxed at a higher rate than other toys. It thus became necessary to sharpen up the definition of ‘doll’ (should Chief Cherokee be taxed at the higher rate?) Accordingly, the Harmonized Tariff schedule defined dolls as being distinguished from toys by “representing only human beings and parts and accessories thereof.”<sup>6</sup> It seems my father was right about Chief Cherokee after all. But of course this makes G.I. Joe a doll too (I don’t know if my father was consistent on this point).

<sup>6</sup> <<http://lawandthemultiverse.com/2011/12/27/are-the-x-men-human-federal-court-says-no>> (last accessed Aug. 2013).

## 10 INTRODUCTION

You might think this settles the matter, but as is often the case, no matter how much we sharpen a definition we run into difficult cases. In this instance the problem arose when attorneys for a company that imported X-men action figures learned that the import tariffs on dolls was significantly higher than the tariffs on mere toys. Since they were paying the doll rate, they went to court to establish that their action figures were not dolls. At issue: whether the X-men or their villains were “human,” a problem complicated by the fact that they are fictional.

The Court of International Trade agreed with the company (Toy Biz, Inc.) and held that the Fantastic Four and related villains, Spider-Man and related villains, etc. were all non-human. (Toy Biz, Inc. v. United States, 248 F.Supp.2d 1234 (Ct. Int’l Trade 2003).) In my view this decision reflected a massive misunderstanding of the relevant comic book characters. Spider-Man and Hulk, despite their mutant ways, are still fundamentally human. But then, no one asked me. Ultimately, the Harmonized Tarriff schedule collapsed the distinction between doll and toy, which was probably a good idea.

### *‘Sandwich’*

One semester while teaching a philosophy of language course I gave my students the assignment of identifying a dispute about meaning that was being played out in the press. I was expecting disputes about ‘person’ or ‘terrorist’ and I got plenty of those, but one student came up with the example of ‘sandwich.’

‘Sandwich’ is an interesting case; there are open-faced sandwiches and wraps and burgers and I suppose croque-monsieurs that one might or might not put in the range of ‘sandwich.’ I was only surprised to learn that the modulation of the definition has legal consequences. In fact, the question came before Pennsylvania Judge Jeffrey Locke in 2006.<sup>7</sup>

An individual franchise in an American restaurant chain/bakery called Panera Bread objected when someone attempted to open a franchise from another restaurant chain—Qdoba Mexican Grill—in the same shopping mall in Shrewsbury, Massachusetts. At issue was the

<sup>7</sup> <<http://www2.courthousenews.com/onpoint/burrito.pdf>> (last accessed Aug. 2013).

fact that Panera's lease with the mall guaranteed that there would be no other sandwich shops in the mall and, argued Panera, a burrito is a kind of sandwich. So it fell to the judge to determine whether a burrito was, in fact, a sandwich. Judge Locke deferred to *Webster's Dictionary* on the following definition of sandwich: "two thin pieces of bread, usually buttered, with a thin layer (as of meat, cheese, or savory mixture) spread between them." The judge then defined a burrito as "typically made with a single tortilla and stuffed with a choice filling of meat, rice, and beans." All good news for Qdoba.

But for author Amanda Hess, it all raised more questions that it answered.

Does an open-faced sandwich constitute a sandwich, despite the lack of sandwiching employed in its construction? If so, is bruschetta a sandwich?

Buttered toast? Pizza?

What if you fold the pizza in half? Must the unifying exterior item be split in two in order to constitute a sandwich? Is a hot dog a sandwich? A submarine roll split in the middle, but with a hinge still hanging on? Is an omelet a sandwich?

A note on methodology: Is it necessary to consume the sandwich with one's own two hands? If one were to douse a sandwich in gravy, would it neutralize the sandwich, converting it into nothing more than a bread-based entree?

If we'll accept a hinge in a sandwich, what about a filling that's encased on two sides? On all sides? Is a kolache a sandwich? A pastry? A corn dog? A calzone? An egg roll? A dumpling? A pop tart? Is a wrap a sandwich?<sup>8</sup>

The idea of a sandwich being encased on all sides had its advocates. Hess spoke with Ian Chillag, who filed reports on his sandwich consumption for a segment on National Public Radio's show *Wait Wait . . . Don't Tell me* called "Sandwich Mondays."

We define sandwich as a 'protein encased in bread product' . . . That way it can include things like the Dunkin' Donuts Pancake Sausage Bites, which is barely even a food, let alone a sandwich. We just figure the more open our definition, the wider the variety of things we can eat and still refer to it as work.

But a quick perusal of the Sandwich Mondays blog suggests a lack of consistency in Chillag's definition. For example on June 11 2012 he

<sup>8</sup> <<http://www.good.is/posts/is-a-burrito-a-sandwich>> (last accessed Aug. 2013).

## 12 INTRODUCTION

blogged about “the pritomas,” a “sandwich devised by a kid,” which consisted of hummus, pickles, and corn chips served “open faced” on an English muffin.<sup>9</sup>

This issue of open-faced sandwiches turns out to have consequences in the bureaucracy of the United States Government. Hess discovered that the United States Department of Agriculture (USDA) had a definition according to which a sandwich “must contain at least 35% cooked meat and no more than 50% bread.” A burrito, on the other hand, according to the USDA, is a “Mexican style sandwich-like product consisting of a flour tortilla.” Why is this relevant to the bureaucracy? Well because the USDA does not regulate sandwiches involving two slices of bread—they only regulate open faced sandwiches. Sandwiches involving two slices of bread are covered by the US Food and Drug Administration (FDA).

Hess offered that perhaps the definition of ‘sandwich’ should not be left up to courts and governmental bureaucracies.

I may not know what a sandwich is, exactly, but I do know that it’s made for everyone. It cannot be defined by courts of law, government directives, or books alone. Its definition must be decided by the people.<sup>10</sup>

This is more or less the position I’ll be defending. It’s not that courts and government bureaucracies can’t or shouldn’t make these decisions; it’s that they are only making the decisions for those within their purview. The decisions don’t have “semantic reach,” which is to say that a government bureaucracy may need to modulate the meaning of ‘sandwich’ for regulatory purposes, but it does not follow that we are compelled to use that modulation. We are speaking different microlanguages.

So far I’ve been discussing humorous cases like ‘doll’ and ‘sandwich’ because they make it vivid that when we talk about what a word means we are not trying to fit the definition to some pre-existing concept—we are not talking about some concept of sandwich or doll that exists

<sup>9</sup> <<http://www.npr.org/blogs/waitwait/2012/06/11/154770514/sandwich-monday-the-pritomas>> (last accessed Aug. 2013).

<sup>10</sup> <<http://www.good.is/posts/is-a-burrito-a-sandwich>> (last accessed Aug. 2013).

independently of the definitions that we supply for these things. Plato's heaven, if it exists, does not have the form of sandwich sitting there. We can offer definitions, and we can modify those definitions, but the value of these definitions ultimately depends upon our interests. This is not to say that there is no right or wrong way to define the term once the interests are fixed. As we will see in Chapter 2, there are clear norms for word meaning litigation.

I also think it would be hasty to dismiss these kinds of lexical disputes (about cases like 'sandwich' or 'doll') as being trivial or a waste of time when they occur outside of the legal or regulatory realm. These discussions may well have value all their own as a kind of lexical grooming. We litigate word meanings with our friends for fun, but in the process perhaps we are constructing a shared language, or at least honing skills that will serve us better when we confront more pressing lexical disputes. A good example of a more pressing case would be how we are to define 'hactivist'.

### *'Hactivist'*

In 2012 an example of lexical warfare unfolded in the treatment of the term 'hactivism', and the dispute over the proper modulation continues to this day. The dispute is interesting in that it incorporates all of the elements of lexical warfare I have just discussed. There had been an ongoing effort to redefine what 'hactivism' meant and what kinds of activities it described, and at the same time there had been an effort to tarnish the label with negative affect so that anyone who chose to label themselves 'hactivist' would do so at their peril.

To a first approximation a hactivist is someone who repurposes technology to effect social change, but there is a conflict between those who want to change the meaning of the word to denote immoral, sinister activities and those who want to defend the broader, more inclusive understanding of 'hactivist'. Attendant to both these efforts is a fight over whether the term 'hactivist' is to have negative or positive affect. Let's start with those who were trying to change the meaning so that it denoted sinister activities.

## 14 INTRODUCTION

In 2012–13 several newspapers and blogs keyed off of Verizon’s 2012 Data Breach Investigation Report,<sup>11</sup> which claimed that in 2011 58 percent of all leaked data was owing to the actions of “ideologically motivated hacktivists.” An example of the concern was an editorial in *Infosecurity Magazine*:<sup>12</sup>

The year 2011 is renowned for being the year that hacktivists out-stole cybercriminals to take top honors according to the Verizon data breach report. Of the 174 million stolen records it tracked in 2011, 100 million were taken by hacktivist groups. Suddenly, things are looking black and white again. Regardless of political motivation or intent, if there are victims of the attacks they perpetrate, then hacktivism has crossed the line. Not OK.

Meanwhile an article in *ThreatPost* proclaimed the following “Anonymous: Hacktivists Steal Most Data in 2011.”<sup>13</sup>

The first thing to note is that both of these media sources were written by and for members of the information security business—it was in their interest to manufacture a threat, for the simple reason that threats meant business for these groups. But is it fair to say that the threat was being “manufactured”? What of the Verizon report that they cited?

The problem is that the headlines and articles, designed to tar hacktivists and make us fear them, did not reflect what the Verizon report actually said. According to the report only 2 percent of the data breaches in the survey were by hacktivists—the bulk of them were by routine cybercriminals, disgruntled employees, and nation states. The “most data” claim stemmed from the fact that precisely two hacktivist actions—both by the now-defunct Anonymous spin-off LulzSec (strictly speaking by the groups Internet Feds and AntiSec) accounted for 58 percent of the data released (these large data dumps stemmed from the actions against HB Gary—a group that went out of its way to pick a fight with Anonymous—and a computer security firm called

<sup>11</sup> <[http://www.verizonbusiness.com/resources/reports/rp\\_data-breach-investigations-report-2012\\_en\\_xg.pdf](http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf)> (last accessed July 2013).

<sup>12</sup> <<http://www.infosecurity-magazine.com/blog/2012/6/7/hacktivism-shades-of-gray-/559.aspx>> (last accessed July 2013).

<sup>13</sup> <[http://threatpost.com/en\\_us/blogs/verizon-hacktivists-steal-most-data-2011-032112](http://threatpost.com/en_us/blogs/verizon-hacktivists-steal-most-data-2011-032112)> (last accessed July 2013).

Stratfor). If you are worried about an intrusion into your system, well then the numbers in the report *actually* suggest it is fifty times more likely that the perpetrator would be a criminal or a nation state or a disgruntled employee than a hacktivist.

In effect, these infosecurity media outlets cited two actions by LulzSec,<sup>14</sup> implicated that actions like this were a principal project of hacktivism, and thereby implicated the imminent threat of hacktivism. Meanwhile, the meaning of ‘hacktivist’ was being narrowed from people who use technology in support of social causes to meaning individuals principally concerned with infiltrating and releasing the data of almost anyone.

Now let’s turn to an attempt to *maintain* the broader understanding of ‘hacktivism’. In the summer of 2012, I went to a birthday party for Daniel Domscheit-Berg, who was turning 34. As it so happens, Daniel had also been the spokesperson for WikiLeaks and, after Julian Assange, the most visible person in WikiLeaks.

The party was to be held in a large house in a small town/village about an hour outside of Berlin. I was expecting to find a bunker full of hackers probing websites with SQL injections and sifting through US State Department cables, but what I found was something else altogether. What I found was an illustration of hacktivism writ large.

When I arrived at the house the first thing I noticed was a large vegetable garden outside. The second thing I noticed was that a tree out front had been fitted out with a colorful knit wool sweater. This was the effort of Daniel’s partner and former Microsoft employee Anka—“Knit hacking,” she called it. And around the small town I saw evidence of her guerrilla knit hacking. The steel polls of nearby street signs had

<sup>14</sup> Strictly speaking, it is sloppy to characterize these hacks as being undertaken by LulzSec. The first hack (of HB Gary) was carried out by a group called Internet Feds. That group subsequently morphed into LulzSec and dissolved within two months. From the ashes of LulzSec, a new group was formed called AntiSec, a member of which (Jeremy Hammond) carried out the Stratfor Hack. It is important to note that the leader of AntiSec was at the time an FBI informant and the hack (and hardware onto which it was downloaded) was under FBI supervision, presumably as an entrapment ploy for Hammond—the point being that were it not for FBI involvement this hack likely would not have happened at all.

## 16 INTRODUCTION

been fitted with woolen sweaters, just like the tree. Most impressively, though, a World War II tank, sitting outside a nearby former Nazi concentration camp for women had also been knit-hacked; the entire barrel of the tank's gun had been fit with a tight colorful wool sweater and adorned with some woolen flowers for good measure. These gestures, I believe, were answers to the attempts to define 'hactivist' as something sinister; they served as ostensive definitions of what hacktivism is and what hacktivists do.

Of course the birthday party had elements of hackerdom understood more narrowly. There were some members of the Chaos Computer Club (a legendary hacker group), and there was a healthy supply of Club Mate—the energy drink of choice of European hackers—but the real story was something else; it was first and foremost about the do it yourself (DIY) aesthetic—planting your own garden, knitting your own sweaters, foraging for mushrooms, and counting on a local friend to bag you some venison. What part of this lifestyle was the hacktivism part? Daniel and his friends would like to say that all of it is.

My point here is that among the things happening was an attempt to defend the traditional, less sinister understanding of 'hacktivism' and perhaps broaden it a bit while adding some positive affect to boot. What they were trying to say is that *hacking* is fundamentally about refusing to let any technology cow us into submission and it is about refusing to be intimidated by any technology. It is about *understanding* the technology and acquiring the power to repurpose it to our individual needs. *Hacktivism*, on their view, was about taking this understanding and power and using it for the good of the many—i.e. to make the world a better place. Moreover, they were saying that a true hacktivist doesn't favor new technology over old—the hacktivist simply refuses to be limited to pre-packaged out-of-the-box technologies. What is critical is that the technologies be in our hands rather than out of our control. This applies to the technologies for food production, technologies for how we shelter and clothe ourselves, and of course the technologies by which we communicate with one another.

What is interesting about this particular episode of lexical warfare was the way it was fought out—with some media outlets of the infosecurity

industry using lexical warfare to create a threat (and more business for themselves) and the hacktivists responding by using gestures of hacking and hacktivism to secure the broader understanding of ‘hacktivism’.

### *‘Journalist’*

In the summer of 2013 the question of how to define ‘hactivist’ gave way to the question of how we should define ‘journalist’. The issue became salient when several United States journalists found themselves under government scrutiny for publishing leaks from whistleblowers and content that had been acquired by hacktivists. The most famous instance of this was the reporter Glenn Greenwald, who assisted the NSA contractor Edward Snowden in his leaking of classified information about NSA surveillance programs.

The issue involving Snowden became salient when the *New York Times* ran an article characterizing Greenwald as an “activist” and “blogger” but withheld the honorific ‘journalist’ despite the fact that he was breaking big stories in the British paper the *Guardian* at the time. Just a few days later, Alexa O’Brien was issuing reports on the trial of Chelsea (then Bradley) Manning, a United States Army private who released millions of pages of secret documents to WikiLeaks. The *New York Times* referred to O’Brien as an “activist” who was “transcribing” the trial, despite the fact that the *Times* was drawing on O’Brien’s work for their own reporting.

O’Brien responded first with an angry letter to the *Times* reporters who wrote the story about her. O’Brien’s letter is worth repeating in its entirety.<sup>15</sup>

Dear Mr. Carr and Mr. Somaiya,

I expect that you will correct your recent article on the U.S. Investigation of WikiLeaks found here:

<https://www.nytimes.com/2013/06/25/world/europe/wikileaks-back-in-news-never-left-us-radar.html?&smid=tw-nytmedia&pagewanted=all>

<sup>15</sup> <[http://www.alexaoobrien.com/secondsight/letter\\_to\\_david\\_carr\\_ravi\\_somaiya\\_and\\_the\\_new\\_york\\_times.html](http://www.alexaoobrien.com/secondsight/letter_to_david_carr_ravi_somaiya_and_the_new_york_times.html)> (last accessed July 2013).

## 18 INTRODUCTION

I am a journalist—and the proper title for me is journalist, most especially because Mr. Somaiya has solicited information published by me in my capacity as a journalist—and I am more than happy to publish my detailed and lengthy email exchange with him for the public.

Mr. Carr, Mr. Somiya, Mr. Bill Keller, The New York Times and other publications have used or linked to my work.

I have been a credentialed member of the press at Fort Meade, MD for 18 month.

My work covering the Manning trial was short listed for the 2013 Martha Gellhorn Prize for Journalism (not activism).

I have received a grant from the Freedom of the Press Foundation for journalism for my coverage of the Manning trial (not for activism).

I find the term activist used here by Mr. Carr and Mr Somaiya—pejorative. So, you will accordingly correct your error immediately.

I am at Fort Meade. Where are you, New York Times?

You are reading my journalistic work, using my journalistic work, capitalizing off of my journalistic work, and linking to my journalistic work about the largest criminal investigation ever into a publisher and its source.

More importantly, you are not here.

Best,

Alexa O'Brien

This letter is interesting in a number of respects. Notice first of all she points out that 'activist' has a pejorative affect—and probably it does among *New York Times* employees and many of its readers. It points to O'Brien's institutional credentials (e.g. a grant from the Freedom of Press Foundation and the Martha Gellhorn Prize for Journalism), she observes that the *Times* has relied heavily on her journalistic product for its articles, and finally she calls out the *Times* for not having representation at one of the most important trials of the century.

The question of whether Alexa O'Brien should be labelled 'journalist' is important not just because of the prestige attaching to the term (which is undeniably important) but because there are recognized protections for individuals who have received the honorific title 'journalist'. In particular, they receive a "qualified privilege" that allows them to withhold the names of their sources. But, as Illinois Senator Dick Durban observed in an Op/Ed piece, this still leaves matters quite open-ended.

In [*Branzburg vs Hayes*], the Supreme Court ruled that there was no absolute privilege for journalists to refuse to reveal sources to a grand jury. The ruling did, however, seem to recognize a qualified privilege for journalists. Today, some federal courts recognize a qualified privilege for journalists, while others do not.

The vagueness of this decision has led 49 states, including Illinois, to recognize a journalist privilege by statute or common law. These laws state that a protected journalist cannot be compelled to disclose sources or documents unless a judge determines there is an extraordinary circumstance or compelling public interest.

But who should be considered to be a journalist?<sup>16</sup>

Durban's question is a good one, with myriad consequences. He gives some of the possible answers (none of which he likely would endorse).

Is each of Twitter's 141 million users in the United States a journalist? How about the 164 million Facebook users? What about bloggers, people posting on Instagram, or users of online message boards like Reddit?<sup>17</sup>

Needless to say, in the age of WikiLeaks and the surveillance state, this is a very important question and there are very high stakes for the players.

Of course questions like this are apt to be questions that journalists (however understood) have to grapple with on a regular basis. Does water boarding fall under the range of 'torture'? Was the sleep deprivation that Manning was forced to endure in the range of 'torture'?

<sup>16</sup> <<http://www.suntimes.com/news/otherviews/20978789-452/sen-dick-durban-its-time-to-say-whos-a-real-reporter.html>> (last accessed Aug. 2013).

<sup>17</sup> <<http://www.suntimes.com/news/otherviews/20978789-452/sen-dick-durban-its-time-to-say-whos-a-real-reporter.html>> (last accessed Aug. 2013).

## 20 INTRODUCTION

The issue extends to questions like what falls in the range of ‘terrorist’ and ‘weapon of mass destruction’. You might be thinking that ‘weapon of mass destruction’ would have a narrow range, including things like nuclear weapons, but like all words and phrases this one is dynamic, and the government, even while it is narrowing the meaning of ‘journalist’, has been broadening the meaning of ‘terrorist’ and ‘weapon of mass destruction’. Needless to say, this is a dangerous combination, particularly since establishment journalists have been acquiescing in these modulations even while independent journalists have resisted them.

Let’s consider the case of ‘terrorist’. Bruce Schneier wrote an article in *The Atlantic* titled “Mission Creep: When Everything is a Weapon of Mass Destruction” and he put his finger on the problem.

One of the assurances I keep hearing about the U.S. government’s spying on American citizens is that it’s only used in cases of terrorism. Terrorism is, of course, an extraordinary crime, and its horrific nature is supposed to justify permitting all sorts of excesses to prevent it. But there’s a problem with this line of reasoning: mission creep. The definitions of “terrorism” and “weapon of mass destruction” are broadening, and these extraordinary powers are being used, and will continue to be used, for crimes other than terrorism.<sup>18</sup>

Schneier had examples to back this up too. One was the case of three anti-nuclear passivists including an 82-year-old nun, who cut through a chain-link fence and entered the Oak Ridge nuclear-weapons-production facility in 2012. “While they were originally arrested on a misdemeanor trespassing charge, the government kept increasing their charges as the facility’s security lapses became more embarrassing. Now the protestors have been convicted of violent crimes of terrorism—and remain in jail.”<sup>19</sup> In another instance, a Tennessee government official claimed that complaining about water quality could be considered an act of terrorism.<sup>20</sup>

<sup>18</sup> <<http://www.theatlantic.com/politics/archive/2013/07/mission-creep-when-everything-is-terrorism/277844>> (last accessed Aug. 2013).

<sup>19</sup> <<http://www.theatlantic.com/politics/archive/2013/07/mission-creep-when-everything-is-terrorism/277844>> (last accessed Aug. 2013).

<sup>20</sup> <[http://www.huffingtonpost.com/2013/06/22/sherwin-smith-tennessee-terrorism\\_n\\_3480930.html](http://www.huffingtonpost.com/2013/06/22/sherwin-smith-tennessee-terrorism_n_3480930.html)> (last accessed July 2013).

Similarly, the government has been attempting to broaden the meaning of ‘weapon of mass destruction’, and it has been expanded to the point where it is ridiculously broad. Here Schneier cites political scientist John Mueller:

As I understand it, not only is a grenade a weapon of mass destruction, but so is a maliciously-designed child’s rocket even if it doesn’t have a warhead . . . All artillery, and virtually every muzzle-loading military long arm for that matter, legally qualifies as a WMD. It does make the bombardment of Ft. Sumter all the more sinister.<sup>21</sup>

### ‘Relevant’

While the government narrows the meaning of journalist and expands the meaning of ‘terrorist’ and ‘weapon of mass destruction’, sometimes these meaning modulations are not sufficient to justify some of their actions—for example, the surveillance of individuals who are not obviously related to terrorism.

According to statute, the US National Security Agency (NSA) can collect records “relevant” to the investigation of terrorism. But what does ‘relevant’ mean here? Not, it turns out, what it means in other contexts. In particular, Robert Litt, General Counsel at the Office of the Director of National Intelligence, argued that the phone records of anyone could be relevant.

As in the grand jury and civil discovery contexts, the concept of ‘relevance’ is broad enough to allow for the collection of information beyond that which ultimately turns out to be important to a terrorist-related investigation. While the scope of the collection at issue here is broader than typically might be acquired through a grand jury subpoena or civil discovery request, the basic principle is similar: *the information is relevant because you need to have the broader set of records in order to identify within them the information that is actually important to a terrorism investigation.* [Emphasis added]<sup>22</sup>

<sup>21</sup> <<http://www.theatlantic.com/politics/archive/2013/07/mission-creep-when-everything-is-terrorism/277844>> (last accessed Aug. 2013).

<sup>22</sup> <<http://www.theatlantic.com/politics/archive/2013/07/mission-creep-when-everything-is-terrorism/277844>> (last accessed Aug. 2013).

## 22 INTRODUCTION

Now this modulation is interesting because it is an understanding of ‘relevant’ that actually doesn’t hold in a typical legal context. Schneier put it this way.

[T]he usefulness of average Americans’ phone records is that it enables the sophisticated analyses performed at NSA to detect patterns that will lead to terrorist activity, rather than—as courts have held outside the national security context—requiring that the information itself be potentially pertinent to an investigation. The essential difference here is that rather than being limited to acquiring the personal information that could relate to the case, this new meaning of “relevance” defines the standard as allowing for the collection of information that will be useful simply by virtue of existing.<sup>23</sup>

While governments and agents often work to broaden or narrow the meaning of terms, sometimes individuals feel compelled to dig in on the meaning of a particular term. This digging in rests on an assumption, which I reject in this book, that meanings are stable (fully fleshed out) things and that appealing to an alleged past or even existing meaning of a term can settle matters. When we engage in lexical warfare, we are interested in the question of how a term *ought* to be defined, not what someone may pronounce it to be.

### ‘Marriage’

As we will see, there are plenty of examples of lexical pronouncement, but cases where religious issues are at stake are where one is most apt to find them. A classic example is the question, which remains contentious for some reason, as to whether couples of the same gender can fall in the range of ‘married’. United States politician and presidential aspirant Rick Santorum was asked by *The Iowa Independent* newspaper why he was opposed to same-sex marriage.<sup>24</sup>

Because it changes the definition of an intrinsic element of society in a way that minimizes what that bond means to society.

<sup>23</sup> <<http://www.theatlantic.com/politics/archive/2013/07/mission-creep-when-everything-is-terrorism/277844>> (last accessed Aug. 2013).

<sup>24</sup> <<http://iowaindependent.com/60602/santorum-qa-marriage-for-gays-threatens-religious-freedom>> (ast accessed July 2013).

Marriage is what marriage is. Marriage was around before government said what it was.

It's like going out and saying, 'That tree is a car.' Well, the tree's not a car. A tree's a tree. Marriage is marriage.

You can say that tree is something other than it is. It can redefine it. But it doesn't change the essential nature of what marriage is.

Marriage is a union between a man and a woman for the purposes of the benefit of both the man and the woman, a natural unitive according to nature, unitive, that is for the purposes of having and rearing children and for the benefit of both the man and the woman involved in that relationship.

It is not only right-wing American politicians who argue this way. When the issue of same-sex marriage found its way into the Canadian court system, Canadian philosopher Adèle Mercier filed an affidavit taking issue with an earlier affidavit *against* same-sex marriage filed by fellow philosopher Robert Stainton. Mercier objected to a claim in paragraph 9 of Stainton's affidavit in which he argued that "It is part of the present meaning of the word 'marriage' in our common tongue that it applies only to male-female conjugal unions. In which case, given the present meaning, it is a necessary truth that same-sex couples cannot marry." In a response, Mercier argued that word meanings just aren't as static as Stainton seemed to be saying.

Even if it were true that the word 'marriage' had referred in the past only to pairs of men and women, that would in no way constitute an argument about the word's meaning, nor an argument that the word 'marriage' cannot refer to pairs other than of men and women... The meanings of all words of all languages, with the exception of personal proper names (which refer all and only to well-defined single objects, i.e. to a person), always stretch beyond their current reference. The word 'Canadians' currently applies to a different group of people than it applied to a hundred years ago, and than it will apply to a hundred years hence.<sup>25</sup>

Strictly speaking, Mercier could be saying that the reference of a term can shift even though the meaning is held constant. But I believe that

<sup>25</sup> Affidavit of Dr Adèle Mercier, Ontario Superior Court of Justice (Divisional Court), between Halpern (et al.) and Canada (Attorney General) et al. and between Metropolitan Community Church of Toronto and Canada (et al.), p. 5. Available online at <[http://www.samesexmarriage.ca/docs/adele\\_mercier.pdf](http://www.samesexmarriage.ca/docs/adele_mercier.pdf)> (last accessed Aug. 2013).

## 24 INTRODUCTION

part of her point is that meanings aren't all that stable either. They shift all the time, we litigate what the meanings ought to be, and digging in on one established meaning is just a way of picking a position and failing to offer arguments for it.

*'Organic'*

As the example of 'hacktivist' showed, we don't need to traffic in the realm of laws and regulations to find consequential cases of lexical warfare. Another example that illustrates the play between meaning underdetermination and the dynamic lexicon is the term 'organic'. 'Organic' of course finds its way into laws, but environmentally conscious consumers sometimes advocate more narrow definitions of 'organic', if only to guide their personal purchases and consumption. A standard definition would be that a food product is organic if is grown without the help of pesticides, but for many, 'organic' is modulated to a much more narrow understanding. For example, one might be concerned with water usage, with whether (in the case of animals) organic implies they are free range, or with whether the crops are planted so as to not exhaust the soil. An article in the *New York Times* identified some of the concerns that might figure into the modulation of 'organic'.<sup>26</sup>

Some organic standard setters are beginning to refine their criteria so that organic products better match their natural ideals. Krav, a major Swedish organic certification program, allows produce grown in greenhouses to carry its "organic" label only if the buildings use at least 80% renewable fuel, for example. And last year the Agriculture Department's National Organic Standards Board revised its rules<sup>27</sup> to require that for an "organic milk" label, cows had to be at least partly fed by grazing in open pastures rather than standing full time in feedlots.

But each decision to narrow the definition of "organic" involves an inevitable tug-of-war among farmers, food producers, supermarkets and environmentalists. While the United States' regulations for organic certification

<sup>26</sup> <<http://www.nytimes.com/2011/12/31/science/earth/questions-about-organic-produce-and-sustainability.html>> (last accessed Aug. 2013).

<sup>27</sup> <<http://www.cornucopia.org/2010/02/new-usda-rules-establish-strong-organic-standards-for-pasture-and-livestock>> (last accessed July 2013).

require that growers use practices that protect water resources, it is hard to define a specific sustainable level of water use for a single farm “because aquifer depletion is the result of many farmers’ overutilizing the resource,” said Miles McEvoy, head of the National Organic Program at the Agriculture Department.

I believe the “tug of war” or lexical warfare in the case of ‘organic’ is particularly interesting because it shows just how wide-ranging the relevant criteria for the modulation might be. Just thinking in terms of pesticide use is not enough; one might also take into account the treatment of animals, the efficient utilization of resources by the farm, and even the very context-sensitive issue of whether the farm is contributing to the depletion of the aquifer.

### 1.3 Unreflective Entrainment

The examples I’ve discussed so far all involve cases where individuals have staked out positions on word meanings and have advocated for those word meanings. In this sense, the way in which they coordinate on word meaning with one another is reflective. This is not always the case. Much of our lexical coordination with our discourse partners is part of a collaborative process that is in many cases automatic and unreflective. Clark (1992) has called this process “entrainment” and it is an excellent if perhaps still metaphorical term for the process.

The original meaning of ‘entrainment’ has to do with the behavior of coupled oscillators in classical physics. There is a great story about the discovery of entrainment by the 17th-century Dutch scientist Huygens, who among his numerous accomplishments invented the pendulum clock. Here is how Huygens described the discovery in a letter to his father penned in 1665.

Being obliged to stay in my room for several days and also occupied in making observations on my two newly made clocks, I have noticed an admirable effect which no one could have ever thought of. It is that these two clocks hanging next to one another separated by one or two feet keep an agreement so exact that the pendulums always oscillate together without variation. After admiring this for a while, I finally figured out that it occurs through a kind of sympathy: mixing up the swings of the pendulums, I have found that within a half

## 26 INTRODUCTION

hour they always return to consonance and remain so constantly afterwards for as long as I let them go. I then separated them, hanging one at the end of the room and the other fifteen feet away, and noticed that in a day there was five seconds difference between them. . . . When in consonance, the pendulums do not oscillate parallel to one another, but instead they approach and separate in opposite directions.<sup>28</sup>

Huygens showed that slight vibrations were being transmitted between the clocks. In one interesting experiment, he hung the clocks on planks that in turn were placed on rickety chairs that were positioned back-to-back and then he put the pendulums out of phase. Initially there was a period of radical shaking, but the system stopped vibrating as the pendulums synchronized (this again took about a half hour).

Huygen's experiment was a great illustration of how even inanimate systems can synchronize. In this case, the vibrations caused by the individual pendulums had effects on the action of the other pendulum up to the point where their effects were mutually reinforcing. They became entrained.<sup>29</sup>

Huygen's clocks were an example of entrainment in a physical system, but it extends to biological systems as well. There has been considerable research on entrainment across a number of areas of science over the past decade, some of it compiled in a popular book by Strogatz entitled *Sync*. One of the key examples from that book involves fireflies in Southeast Asia that flash in unison. When this natural phenomenon was initially discovered the explanations ranged from the idea that there must be a boss firefly that they all followed to flat out denial of the facts.

For example in 1917, an author in the journal *Science* remarked, "some twenty years ago I saw, or thought I saw, a synchronal or simultaneous flashing of fireflies. I could hardly believe my eyes, for such a thing to occur among insects is certainly contrary to all natural laws." In 1918, George Hudson wrote that "if it is desired to get a body of men to sing or play together in perfect rhythm they not only must have a

<sup>28</sup> Translation from Strogatz (2003).

<sup>29</sup> Interestingly, the theory of coupled oscillators has been applied in metrical phonology. See Barbosa (2002) and O'Dell et al. (1999).

leader but must be trained to follow such a leader.” Eventually, however, an alternative explanation emerged. Strogatz put the explanation as follows.

Taken together, the two clues suggested that the flash rhythm was regulated by an internal, resettable oscillator. And that immediately suggested a possible synchronization mechanism: In a congregation of flashing fireflies, everyone one is continually sending and receiving signals, shifting the rhythms of others and being shifted by them in turn. Out of the hubbub, sync somehow emerges spontaneously.

Thus we are led to entertain an explanation that seemed unthinkable just a few decades ago—the fireflies organize themselves. No maestro is required and it doesn’t matter what the weather is like. Sync occurs through mutual cuing, in the same way an orchestra can keep perfect time without a conductor. What’s counterintuitive here is that the insects don’t need to be intelligent. They have all the ingredients they need. Each firefly contains an oscillator, a little metronome, whose timing adjusts automatically in response to the flashes of others. That’s it.

Strogatz goes on to argue that the phenomenon is quite widespread in nature and, because it is driven by low-level mathematical and physical properties, the synchronization is inevitable if the initial conditions are right. Indeed, as Strogatz puts it, “the tendency to synchronize is one of the most pervasive drives in the universe, extending from atoms to animals, from people to planets.”

But what about lexical synchronization? Presumably, the talk of lexical entrainment is metaphorical because we don’t think about meanings oscillating (although it would be interesting to try and make sense of the idea of meanings oscillating between alternatives in a semantic space of some form). On the other hand it does make sense to think that when we are in a state of meaning mismatch with our collaborators it generates perturbations—misunderstandings, confusion, and of course not a little expenditure of cognitive labor to right things. It would make sense for us to be optimized for synchronizing, but how would this work?

There are actually two questions to be answered. First, how does the unreflective synchronizing take place and what are the mechanisms by which it comes into effect? Second, what role does semantic deference play in this and is it a push-me (imposing your will) or pull-you (copying the person in power) strategy?

## 28 INTRODUCTION

Let's begin with the process of synchronization. Here it might be useful to look at work that has been done in the area of Conversation Analysis (CA), for example Sacks (1995), Sacks et al. (1974), Sidnell (2010). While this work comes larded with a fair bit of anti-Chomsky ideology and has come under criticism from Searle (1987), on the grounds that its rules for turn-taking are implausible rules,<sup>30</sup> we can ignore that and get to the data itself, which nicely illustrates ways in which people engage in turn-taking (let's set aside why) and in which they modify their linguistic behavior on the fly. While the CA analyses do not always look at the way in which word meanings are modulated (it is more focused on the mechanics of turn-taking), there is plenty of data that can provide us some insights into the process.

The interesting thing about the CA data is how nicely it illustrates that our conversations are not the cleanly scripted exchanges we see in movies, but they typically involve rapid turn-taking (occasionally with overlap), ample repair and self-repair, challenges, and requests for clarification.

Let's look at some examples to get a taste of how the process can work. Our first case (from Schegloff 2007) involves a 14-year-old girl named Virginia introducing a new term to her mother. It is not a new coinage, but the introduction of a term learned elsewhere, now for the benefit of her mom. Her brother's fiancé Prudence asks what it means, and her mom picks up on her daughter's term, but signals that she still isn't clear on what it means.<sup>31</sup>

- 01 Mom: I don't think that  
 02           you should be going to the parties that Beth goes  
                   to. She is  
 03           eighteen years old. An' you are fourteen da[r]ling

<sup>30</sup> I actually believe that Searle's criticism misses the mark, since it trades on his view that rule following must be conscious in principle. For reasons given in Ludlow (2010: ch. 4) I don't agree with this.

<sup>31</sup> I've left out some of the notation as it is not critical for us, but some remains. Underlining indicates stress. Capital letters indicate volume. The material lined up by left square brackets (as in lines 3 and 4) indicate overlap. Numbers in parentheses (as in line 12) indicate length of pause in seconds. In some instances there is cross-talk or noise so that no speaker is identified.

- 04 Vir: [I KNOW, BUT  
ALL THE
- 05 REST OF THE PEOPLE MY AGE ARE GWAFFS.  
I promise. they
- 06 are si[ck
- 07 Mom: [They're what?
- 08 ()
- 09 Vir: GWAFFS
- 10 ??? ()
- 11 Pru: What's a gwaff?
- 12 (3.1)
- 13 Vir: Gwaff is jus' someb'dy who's really... I just- ehh!
- 14 s- immature. You don't wanna hang around with  
people like that
- 15 (1.9)
- 16 Mom: Well, don'tchyou think thet thuh eighteen year olds  
an thuh
- 17 twenny year olds think you're a gwaff?
- 18 (0.8)
- 19 Mom: Whatever a gwaff might be?
- 21<sup>32</sup> Pru: [ehh huh!
- 22 Vir: [eWell not if I date 'em, I mean gosh!
- 23 Pru: ehh!

Notice that Virginia seems know that her mom won't know what a gwaff is, and she provides a clue by saying, without being prompted, that "they are sick." When her mom uses the term 'gwaff' but then adds "whatever that may be" she is, per the norm, deferring to her daughter's usage, but also signaling that she is only temporarily taking it on. But clearly she understands it well enough to challenge her daughter—if a gwaff is someone who is immature, then don't the older kids think she is a gwaff? But now we are sorting out what 'immature' means.

<sup>32</sup> Line numeration in the original skips from 19 to 21.

## 30 INTRODUCTION

A bit later in the conversation her brother Wesley interjects another concern at line 53: What if the older kids are taking advantage of her? Now it is mom's turn (line 61) to ask what 'taking advantage of' means.

- 46 Wes: what- all these young people yer own age. You don't like  
 47 tuh do thuh same things they do?  
 48 (0.9)
- 49 Vir: No I hang around [some people my age but they hang  
 around ...
- 50 Wes: [(That's enough.)  
 51 ... older people  
 52 (2.0)
- 53 Wes: You're not worried about 'um takin' advantage of yuh?  
 54 (1.5)
- 55 Vir: W[ho  
 56 ??? [someone sneezes  
 57 *Nuh-(h)O!*  
 58 ??? (huh huh 'hh)  
 59 (3.5)
- 60 Vir: Thu only time any[body  
 61 Mom: [Whaddy mean by that
- 62 Pru: Mm hm hm!  
 63 (0.6)  
 64 Wes Wull 'ey just- the[y'll say thin]gs and ... they'll ...
- 65 Pru: [( )]  
 66 Wes ... lie to yuh, 'n you won't know when they're tellin'  
 you thuh  
 67 truth  
 68 (.)
- 69 Vir: Buh yes I will  
 70 (2.0)
- 71 Wes: Whatever.

Sometimes the term is completely familiar, and people may even be on the same page, but we need mechanisms that allow us to check and make sure from time to time; we need what communications scientists

call “error correction code.” To illustrate, consider the word ‘flash’, which on one reading means to expose yourself briefly. But to what degree? And for how long? Clearly we have a case of meaning under-determination on our hands. In the following exchange, transcribed from an episode of Donald Trump’s reality TV show, *The Apprentice* (from Sidnell 2010), we get an example of how semantic authority asserts itself and how it checks semantic conformity in the course of conversation.

The Apprentice: Boardroom

- 01 Trump: Ivana. You flashed (0.4) a group of people  
 03 Ivana: look [this ]  
 04 Trump: [no no no] did that happen?  
 05 Ivana: it happened? but it happened for a reason.  
 06 Trump: why  
 07 Ivana: because I knew—okay we had gone through  
 08 a lot of product we [only had]  
 09 Trump: [what does] flash mean  
 10 you ripped down your pants? [what does that mean]  
 11 Ivana: [I was wearing - ]  
 12 I was wearing a bikini  
 13 (0.4)  
 14 an- an let’s not blow this out of proportion,  
 15 I was wearing bikini shorts I wear  
 16 Caroline: We [haven’t] said anything yet so [relax ]  
 17 Ivana: [more ] [I know I know]  
 18 I’m really just defensive a[bout this because ]  
 19 Trump: [go ahead I’d like to hear at]  
 20 Ivana: Um  
 21 Trump: But you did flash  
 22 Ivana: I did but it was a gimmick it was a gimmick  
 23 just like [girls ]  
 24 Trump: [Did it work?]  
 25 (0.3)

## 32 INTRODUCTION

- 26 Ivana:           It did work [I sold a candy bar for twenty bucks     ]  
 27 Trump:                   [Oh=rilly, but you're on the losing team]  
 28                   (0.8)  
 29 George:       You sold a candy bar for twenty dollars?

It is interesting that Trump, who introduces the term ‘flashed’ (in line 1), asks what it means (in line 9), but pretty clearly he isn’t asking so as to find out. He is asking because he wants to make sure that Ivana is on the same page as him—that what she did (the details seem irrelevant to Trump) was a clear case of flashing. Notice that Ivana could have protested that what she did was not a case of flashing, but she agreed for purposes of the discussion to concede that what she did was clearly in the range of ‘flashing’.

There are other attempts at semantic policing going on in this dialogue. At line 24 Trump asks Ivana if her strategy “worked.” One might wonder what he meant by this. Ivana takes ‘work’ to mean it allowed her to successfully sell candy bars, but Trump is having none of that explanation, saying “oh really, but you’re on the losing team.” Trump’s point seems to be that it only worked if you won. Notice also that there is an interesting challenge to the contention that she was selling a candy bar. I’ve left out some of the phonetic details of this transcript but they show that George, by stressing and lengthening the first syllable of ‘candy’ in line 29, seemed to be challenging the assumption that candy bars were what she was selling. This too might be thought of as a metalinguistic correction—what sort of activities can count as in the range of ‘selling a candy bar’. Perhaps it is not enough that someone gets money and someone else gets a candy bar. Trump goes on to agree with George. Ivana wasn’t selling a candy bar, but the flash.

In a bit we will get to the matter of deference, but notice first that this isn’t just passive semantic deference—it is Trump imposing his semantic authority. In that context (it’s his show after all and she is on the hot seat) Ivana really has no choice but to defer to Trump on linguistic usage. She can offer up facts of the matter, but Trump seems to get the final say on the appropriate modulation of ‘flash’, ‘worked’, and ‘selling a candy bar’. This is not surprising given the enormous power imbalance. (Ivana was, by the way, fired.)

Next consider a case where it appears Jim is asking for an error correction check by asking what Roger meant by ‘just agreeing’ and notice how two separate interlocutors collaborate on cueing Jim in on the correct modulation (Sidnell 2010, originally from Sacks et al. 1974).

- 01 Roger: Are you just agreeing because you feel you wanna uh  
 02 Jim: Hm?  
 03 Roger: You just agreeing?  
 04 (0.4)  
 05 Jim: What the hell’s that.  
 06 Al: It’s Agree[ing]?  
 07 Roger: [Agreeing.  
 08 Jim: Agreeing  
 09 Roger: Yeah  
 10 (:): A[greē’n  
 11 Al: [With us. Just going along with us  
 12 Jim: No.  
 13 (0.4)  
 14 Roger: Saying ‘yes, yes’ [hehheh hh hehhh hh hehheh hh  
 15 Jim: [well i-i-it’s, it’s true.  
 16 Everything he said is true, so

It’s not clear that Jim was really looking for a definition (he might have been asking “why the hell are you asking that?”) but a definition of ‘agreeing’ is provided, and it seems like there is also some word meaning modulation here. ‘Agree’ could be taken to mean to concur that something is true, but that isn’t how Al and Roger are using the term here. When Al says in line 11 “Just going along with us” it is clear that by ‘just agreeing’ he meant he was asking if Jim was saying ‘yes’ even if he didn’t believe it was true (this seems to be Roger’s point as well, or rather this was Roger’s recasting of Jim’s definition).<sup>33</sup>

<sup>33</sup> Sidnell (2010) seems to think the speaker Jim has misunderstood the word as ‘green’ but I don’t see any evidence for this. In any case Al and Roger take him to be offering a serious question and they respond in kind, so we still get a good sense of how these meaning clarifications can work. This particular fragment is also discussed in Sacks (1995) and Sacks et al. (1974).

## 34 INTRODUCTION

Let's consider a final example (again from Sidnell 2010) that turns on how we are to understand the expression 'dressed up'. Does it mean that the clothes are dress clothes or does it also mean that the clothes should be in good condition?

- 01 Kathy: You got all dressed up? just to see us?  
 02 Reuben: Are you kidding?  
 03 (pause)  
 04 Frieda: I'm all ripped.  
 05 Kathy: Oh yeah  
 06 Frieda: Yeah  
 07 Kathy: I can see the hole  
 08 Frieda: all over  
 09 Reuben: Don't you recognize my uniform?  
 10 Kathy: Yes. No, I meant Frieda was wearing a fancy dress.

In this case Frieda moves to narrow the meaning of 'dressed up' so that it doesn't include clothes with holes in them. I suppose it is plausible that Kathy was working with the narrow definition initially, but I suspect that she is just going along with Frieda on this. Why bother arguing about what 'dressed up' means? Especially in this case. Frieda is pushing for the narrow modulation because she is looking for a face-saving way out of the social embarrassment of being overdressed. Kathy is happy to accommodate her, although in other circumstances where being dressed up was called for she would presumably insist that Frieda *was* dressed up. No one is explicitly talking about word definitions in the conversation, but they definitely have modulated the meaning of 'dressed up'. The shift in definition happens below the surface and it is part of an attempt to smooth out a socially awkward situation. Reuben, meanwhile, seems to be oblivious to everything happening.<sup>34</sup>

<sup>34</sup> It's interesting how not clued in he is. It is almost as though Kathy and Frieda are having a private conversation—as though he isn't a participant in the microlanguage at all.

Now, obviously there is plenty of subtle social reasoning going on here, and one ultimately wants a theory of how that works, but my point here is really about how the meaning of ‘dressed up’ is just a bit player in all of this, and it isn’t all that reflective. A modulation is offered and it is more or less unreflectively taken on. It is like a game piece that can be moved about to satisfy other concerns. When words are modulated we generally play along, even if our interlocutor isn’t a billionaire like Trump.

Work on conversational analysis suggests that we often work together in this way. One sees it very clearly in the introduction of names for things; likewise for pronunciation. Someone introduces a term or pronunciation X, the interlocutor uses an alternative Y, and then the initial speaker complies by using Y. Presumably the change to Y is conceded because there would be no point in *offering* the repair unless there was a reason (repair has a cost—the least effort principle would say that all other things being equal we should go with the flow). But are we really this deferential? Often we are.

There many cases where we blindly or at least indifferently adopt the linguistic practices of those around us, apparently for no reason at all. Well, *maybe* it is for no reason.

It is certainly the case that human agents are quite adept at simply doing as their neighbors do. Joshua Epstein, an economist at the Brookings Institution, has shown that one can successfully model group political behavior with a population of cellular automata that basically just do what their neighbors do as long as no new agent comes along and violates conventions.<sup>35</sup>

It is interesting to reflect on whether this behavior, hardwired or not, could count as being rational or normative in some sense. Surely *some* unreflective imitation must be warranted. It would certainly make for an interesting time if all conformity required pause for reflection. Quite apart from making driving an adventure (because of having to

<sup>35</sup> See Epstein (2001).

## 36 INTRODUCTION

reflect on whether driving on the right/left side of the road is the thing to do), many of us would simply be paralyzed with indecision.

In recent years a number of philosophers have pursued the idea that we are entitled to quite a bit of knowledge, including knowledge that we gain from the testimony of others and sometimes just from looking without much serious reflection (see e.g. Burge 2003). One can likewise imagine a similar theory that establishes our semantic warrant for reflexively following our neighbors when they introduce novel lexical items or when they offer modulations of those already in use.

There are moments, however, when our preferred modulations of a word meaning collide or where we have to choose between conflicting modulations. In some cases, as noted earlier, we defer to a perceived semantic authority. In other cases, we actually resist someone's modulation and litigate for our preferred modulation. We will get to the latter type of case in the next chapter. For now I want to stay focused on the issue of deference, and ask the question: just how do we determine who is semantically deference-worthy? Who do we copy? By quantitative economic measures, Donald Trump is successful, but does that make him deference-worthy?

It is one thing to say that semantic deference takes place and quite another to explain how it works. Friend and Ludlow (2004) considered the thesis that deference-worthiness is earned discursively via a series of challenges. This involved a two-level process—first determining whether the interlocutor has salient domain expertise, and second determining whether the expertise has semantic reach in this context. More precisely, we argued that expertise in a domain must be established via a series of interactive “partial knowledge proofs.” The phrase ‘partial knowledge proof’ is a riff on the notion of “zero knowledge proofs” in computer science (in particular in the field of public key cryptography). The basic idea of a *partial* knowledge proof is this: If I have a particular expertise, how can I prove to you that I have that expertise when it is something that you lack? To illustrate the idea, imagine a situation where we are hiring a philosopher in ancient philosophy but no one in the department is an expert in that area. We all

have *some* knowledge of ancient philosophy, of course, but we are hiring in the area because we recognize we are not experts. We resolve this dilemma by issuing a series of challenges to the job candidate. With each question/answer exchange we learn more, allowing our colleagues to press on with deeper and more informed questions. In the end, via this interactive inductive proof procedure, we satisfy ourselves that the candidate is worthy. Or not.

Stacie Friend and I argued that this kind of procedure is more common than one might think, applying even in cases like the meaning of the word ‘cool’ (in the social not the thermodynamic sense). Think about the social dynamics depicted in the 1970s television show *Happy Days*. We might think that Richie and Pottsie always blindly defer to Fonzie on the meaning of ‘cool’, but in fact there are times when challenges are issued, and there are at least person-internal debates about whether Fonzie is really the appropriate arbiter of the extension of the term. Fonzie’s deference-worthiness is constantly subject to challenge, and may well be undermined as we encounter other arbiters of ‘cool’ (as when Richie goes to college) or aspects of Fonzie’s behavior (as when he goes water skiing and jumps a penned-up shark—definitely not cool).

It is an interesting question as to what counts in a decision to defer to Fonzie on the meaning of ‘cool’. Presumably Richie and Pottsie had partial knowledge of the concept, and their deference is not tied to credentials possessed by Fonzie; Fonzie did not have a diploma from the College of Cool. In other cases, however, semantic deference *does* appear to be tied to credentials.

For example, one day a “tree guy” came to my house and, while pruning some trees, identified the trees in my yard. Along the way he assured me he had gone to horticulture school. Did that provide him with the expertise to say which is a beech and which is an elm? Should I defer to him? Well, I’m not much hung up on the question, so I was perfectly happy to adopt his usage. For similar reasons I’m happy to defer to the doctor when she says I can’t have arthritis in my thigh. But *why* do I defer?

## 38 INTRODUCTION

Well, presumably it is not because these experts have pointy heads or impressive accents—it is because the credentials they hold (diplomas, for example) show they have been vetted by a kind of process not so different from the one we used to hire our ancient philosopher—as students they were subject to an interactive inductive proof procedure which convinced their institutions that they had the relevant domain knowledge. It would be interesting to explore this process in more detail, though when we turn to the semantics of word meaning a more pressing question arises: Why does your domain expertise matter *here*?

The point of my question is that, once domain expertise is established, the “semantic reach” of the domain expertise must also be established (e.g. should I defer to the materials scientist when she says that the glass in that window falls under the extension of ‘liquid’ in our conversation? Or is the materials scientist overreaching her jurisdiction when she asks us to adopt her linguistic usage?). In Ludlow and Friend (2004), we considered the idea that this semantic reach can also be established discursively, via a series of challenges. In effect we can think of these as being cases where we challenge someone’s semantic authority, or in any case challenge them on a particular modulation. And of course, this will happen even if there is an imbalance in power relations—necessarily so. Those in a position of semantic authority in a given context are always subject to challenge.

## 2

# Norms of Word Meaning Litigation

In this chapter I go into detail on a handful of cases in which we are consciously aware of disputes about word meaning and in which we litigate or argue about the best way of modulating the term in dispute. I reject the idea that this is just a matter of imposing our will on our interlocutors. Recent work in the theory of argumentation has shed considerable light on this process,<sup>1</sup> but we will need to refit that work for the kinds of considerations we are engaged with here.

I'll begin this chapter with a general description of how we come to notice that there are conflicts in meaning and how we structure the meaning litigation once the conflicts are recognized. I'll then take up an example case that is relatively less controversial—the definition of 'planet'—and use it to construct a model for how meaning litigation works. I'll then turn to more contentious and substantial issues—the definition of 'rape' and the definition of 'person' and begin exploring how disputes about the meanings of those terms can be normative and fail to be normative.

### 2.1 Recognizing and Engaging Meaning Mismatch

When we engage with others in conversation, all of the participants come to the table with a robust linguistic background already in place.

<sup>1</sup> See Ajdukiewicz (1974), Eemeren and Grootendorst (1992), Eemeren et al. (1996), Eemeren (2002) and Budzynska (2006) for examples.

We could say that the words the participants are deploying have residual value from previous entrainments. In many cases, we will have the expectation that we and our discourse partners are already entrained with each other, and assume that we lexically converge with respect to the words we are using—that is, we converge well enough to allow us to engage in conversation. Sometimes we know there is a lack of convergence, for example in obvious cases like when we are in a country where we do not speak the dominant language, and more subtle cases as when we know that our conversational partners have not taken a philosophy class and thus won't take 'pragmatist' to be speaking of a school of American philosophy, or understand that we intend 'valid' to have a very specific meaning.

While sometimes we can see the different takes on word meaning coming, sometimes we do not see this until we are already engaged in conversation. We can call the cues that allow us to recognize semantic mismatch "triggers." These triggers can take different forms.

Often, a few minutes into a conversation, we recognize that we are using a term differently than our communicative partner is. Sometimes we recognize that these differences are differences in modulation. For example, you may have a much broader modulation of 'athlete' than I do, so that for you it takes in racehorses and chess players.

Other times, we may recognize that we need to sharpen a word meaning if we are to successfully resolve some problem or make a decision and take action on it. For example, we might recognize from technological advances that our definition of 'death' needs to be sharpened, so we engage in a discussion about the best way to sharpen it.

Of course the real point of interest is in what happens once we recognize these differences in meaning and we begin litigating them. Let's set aside cases where there is a power imbalance and one participant simply defers to another and let's also set aside cases where we simply agree to disagree; let's consider cases where all sides want to make their case and persuade the other to follow them. Is there a best way to proceed? Or is it simply a matter of who can be the most persuasive?

Earlier I alluded to the theory of argumentation—a theory that is not concerned so much with the form of arguments themselves, as

with the *process* of argumentation and the methods of argumentation that are more apt to yield the correct result. Roughly speaking, the strategies involve an attempt to find beliefs that the discourse partners share, and then reason from those shared beliefs in an attempt to get their discourse partner to defect.

Of course, that doesn't say much about how the reasoning process works, and I think that it is still an open question as to the strategies humans use in this regard as well as an open question as to what strategies are normatively correct—as we will see there is not always an easy way to settle this latter question. What I propose to do in the remainder of this chapter is to examine three cases where word meanings have been litigated in an attempt to illuminate at least some features of the process, and then we will try to get clear on the processes that are in some sense more reliable. I'll begin with a case that is less politically and emotionally charged—the word 'planet'—and proceed to more contentious cases like 'rape' and 'person'.

## 2.2 'Planet'

As most people know, the word 'planet' originally had the meaning "wanderer" and it was used to speak of the celestial objects that did not have a fixed position with respect to the other stars, but moved among them. The original six "planets" were thus Mercury, Venus, Mars, Jupiter, Saturn, The Sun, and The Moon. Subsequent empirical discoveries shook up this taxonomy. We now count The Earth as in the range of 'planet' and The Sun and The Moon as not in the range.

Obviously a definitional shift didn't *need* to happen. We could have held the range of 'planet' constant for reasons of historical continuity and deployed a new term for Mercury, Venus, The Earth, Mars, Jupiter, and Saturn—"large solar orbitals" for example. But we didn't. So why did the definition of planet shift? Why did we modulate the meaning of 'planet' in the wake of scientific discovery?

One possible story is that we took the term 'planet' to pick out things that had a uniform class of properties. When it turned out that not all

of our canonical exemplars within the range of ‘planet’ had these properties, we modulated the word meaning so that it preserved the large subset of cases that happened to share the property of being spherical objects in primary orbit around the sun. If this is right, then we can say that people shifted the definition in such a way as to reflect some important shared property (or properties) of the bulk of the original canonical exemplars in the range of ‘planet’.

More recent scientific discoveries have again called into question the proper definition of the word ‘planet’. To some extent these definitional disputes have made it into the public sphere with the question of whether Pluto should count as being in the range of ‘planet’.

As a bit of background to this case, it is important to understand that Pluto has always been a bit weird as far as planets go. For starters, it is on a different orbital plane than the other planets, and we have recently come to understand that its material composition is rather different than the other planets. Unlike rocky planets like the Earth and Mars, and unlike the gaseous planets like Jupiter and Saturn, it is basically a ball of ice.

What initiated the rethinking of the status of Pluto was a series of discoveries that began in 1992 when Jane Luu and David Jewitt discovered the first Kuiper Belt object. Since then, thousands of additional Kuiper Belt objects have been discovered, several of which are nearly as large as Pluto and some, like Eris, larger.

What pushed the case of Pluto into public consciousness was the demotion of Pluto on February 19, 2000 by Neil Tyson, who was director of the Hayden Planetarium at the American Museum of Natural History. On that day, visitors to the planetarium no longer found Pluto listed among the planets, and instead found this statement of the taxonomy of objects in the solar system.

Five classes of objects orbit our Sun. The inner terrestrial *planets* are separated from the outer *gas giant planets* by the *asteroid belt*. Beyond the outer planets is the *Kuiper Belt* of comets, a disk of small icy worlds including Pluto. Much more distant, reaching a thousand times farther than Pluto, lives the *Oort Cloud* of comets.

On January 22, 2001, the *New York Times* objected to this new taxonomy:

Quietly, and apparently uniquely among major scientific institutions, the American Museum of Natural History cast Pluto out of the pantheon of planets when it opened the Rose Center last February. . . . the move is surprising, because the museum appears to have unilaterally demoted Pluto, reassigning it as one of more than 300 icy bodies orbiting beyond Neptune, in a region called the Kuiper Belt.

Members of the scientific community also weighed in, including Alan Stern at the SW Research Institute: “They [the Hayden Planetarium] are a minority viewpoint . . . It’s absurd. The astronomical community has settled this issue. There is no issue” (quoted in Tyson 2009: 82). As Phil Plait of Sonoma State observed, the dispute had been brought to a head by the fact that there was no extant definition of planet—it seems we had been working with an ostensive definition based on canonical cases.

At the heart of the debate is our very definition of the word ‘planet’. Currently, there isn’t one. The International Astronomical Union (IAU), a worldwide body of astronomers, is the official keeper of names. It has no strict definition of planet, but has decreed that there are nine major planets, including Pluto. This, however, is not very satisfying. If the IAU doesn’t really know what a planet is, how can it know there are nine? (Quoted in Tyson 2009: 104)

The Planetary Definition Committee of the International Astronomical Union subsequently met on August 16, 2006 in an attempt to fill this lacuna. After discussion they proposed a definition that had two components: A planet is an object that

1. Is in orbit around a star but not around another planet.
2. Is large enough for gravity to form it into a sphere but not so large as to cause as to trigger fusion. (As aficionados put it, not so large as to cause deuterium burning.)

But just days later, on August 24, 2006 the general assembly of the IAU rejected this definition and added a third criterion. In addition to the criteria offered by the Planetary Definition Committee, they stipulated that

## 44 NORMS OF WORD MEANING LITIGATION

## 3. The round object has cleared its orbit of debris.

Let's pause and make some observations at this point. First, note that the meaning of 'planet' was underdetermined, even though we may not have recognized it previous to recent scientific discoveries. Those discoveries provided the triggers for us to recognize that we were encountering cases (e.g. some Kuiper Belt objects) that were not determinably in or out of the range of 'planet'. This precipitated the attempt to modulate the meaning of the term 'planet'. The modulation of the meaning was justified by an argument to the effect that, if Pluto is recognized as in the range of 'planet' then many other objects must be as well, and involved an attempt to screen out objects that were large enough to be round, but still sitting in the Kuiper Belt. The resulting modulation added some sharpness (at least in the context of the current state of our solar system) but it was still not completely precise, as noted by Tyson (2009: 118).

The [third] criterion is subtle because without a quantitative account of a clean orbit the criterion can be arbitrarily invoked... Earth continues to plow through hundreds of tons of meteoroids every day... So have we cleaned our orbit? Clearly not... The objective is to assess the total mass of cleanable debris and compare it with the mass of the planet in question. If the debris does not amount to much, then you can claim to have cleaned or dominated your orbit.

It is important to distinguish *narrowing* word meanings from *sharpening* word meanings. Often we can narrow a word meaning without sharpening it. For example, we can stipulate that I'm not bald, but this doesn't sharpen the meaning of 'bald' because it doesn't tell us where the edges are—only that 'bald' does not apply to people with more hair than me. It doesn't fix the edge for people with less hair. Likewise the three-part definition of 'planet' *narrowed* the meaning by excluding cases in our solar system, but it did not *sharpen* it because there was no real attempt to sharpen the notion of a clean orbit. Or more accurately, the definition is sharp enough for our solar system, but not for many others where orbits are in the process of being cleared.

We can also sharpen the meaning without narrowing it. Some definitions of 'planet' preserve Pluto as a planet and are quite sharp. For

example the definition provided by the planetary definition committee was just such a case in that it sharpened up the definition to admit objects that are large enough to be round but not so large as to allow fusion. This admits Pluto and it also gives rise to fewer difficult cases in other solar systems with varying degrees of debris-filled orbits. Of course even this definition is not completely sharp; some objects are large enough for minimal fusion to take place but not enough to credibly be called ‘stars’.

This was recognized by George Wetherill, a planetary scientist at the Carnegie Institute of Washington, who noted that the deuterium-burning definition is precise enough for current cases and can be sharpened if needed.

Distinguish between a planet and a star by deuterium burning. There will be borderline cases, but so what? Some day, when we understand formulation of these bodies much better, finer distinctions can be made... (Quoted in Weintraub 2007: 229)

Our discussion so far as been brief, but it gives us enough resources to begin fleshing out the way word meanings are adjusted and litigated in cases like this. We can enumerate them for consideration.

- (i) Take undisputed cases and argue analogically for new cases (or against familiar cases).

For example, we can argue analogically from traditional planets to inclusion of The Earth (and exclusion of The Sun). Similarly we can argue analogically for or against Pluto.

Reasoning analogically about Pluto, we can say that it is *like* undisputed planets, in that it is (1) not massive enough for fusion, (2) massive enough to form a ball, (3) orbits the sun. It is *unlike* undisputed planets, in that it is (1) mostly made of ice, (2) not on the same plane as the undisputed planets, (3) hasn’t cleared its orbit. It is *more* like other Kuiper Belt objects.

We can also discern some additional principles.

- (ii) Modulations should respect the bulk of canonical cases.

For example, in the original shift in meaning of ‘planet’ it seems that there was an attempt to keep the originally ostended objects within the range of ‘planet’. Obviously we needed to give up some—The Sun and The Moon—and we added one (The Earth) as science advanced.

Some of the debates around the modulation of ‘planet’ held that not only should the canonical cases be preserved but also that they should be in some sense safe. Consider the following passage from Weintraub (2007: 203).

Can we assert that a planet must have a moon? . . . This “must have a moon” requirement would drop Mercury and Venus from the list of planets and make Mars questionable. Did Mars become a planet only after it captured its moons from the asteroid belt a few hundred million years ago?

My interest here is with Weintraub’s discussion of Mars. Moons come and go—they can be captured or go wandering off under the right conditions (another planet pulling them out of orbit, for example). Weintraub seems to be suggesting that definitions should not only preserve canonical examples, but also that they should not be subject to contingencies like the coming and goings of their moons.

(iii) Modulations should track (not cross-cut) important properties.

What makes this criterion interesting is that often the properties are understood to be important in the wake of scientific discovery. For example, when we discovered that the things we are calling ‘planets’ are not wandering stars, but most of them are bodies in primary orbit around the Sun, we took this to be the crucial property, and so we adjusted the definition to respect and not cross-cut this property. Thus the meaning of ‘planet’ was modulated so that the Sun and Moon were excluded from the range of ‘planet’ and so that Earth was included.

This is somewhat similar to the case of ‘Polio’. It came as a discovery that not everything diagnosed as “Polio” was caused by the newly discovered virus, so if we wanted the term ‘Polio’ to track important properties and not cross-cut them then many of the conditions formerly diagnosed as being Polio would have to be modulated out—they would be classified as something else, caused by “Non-Polio Enteroviruses” (NPEVs).

Michael A'Hearn, a professor of astronomy at the University of Maryland, puts the scientific motivation for this as follows:

Why do we, as scientists, care how Pluto (or anything else) is classified? ... Scientists put things into groups, the members of which share common properties, in order to find patterns that will enable us to better understand how the bodies work or how they became what they are. If we are interested in origins, then it is clear with our present understanding (which might change in the future) that free-floating bodies of mass comparable to Jupiter are not in the same class as Jupiter itself. Similarly, it is clear that Pluto is not a planet like Jupiter but is rather a planet like the numerous Plutinos that live in the 3-2 libration with Neptune. Thus Pluto should be classified as the largest Plutino. (Quoted in Weintraub 2007; 229)

Tyson (2009: 77) offered a similar justification.

We looked across the solar system and asked ourselves what physical features about planets and other objects could be taken together and discussed as common properties of phenomena, allowing us to compare and contrast objects in whatever way those families would naturally delineate. ... Pluto was displayed with other Kuiper belt objects but we neither counted these objects nor made a list of who is or is not a planet.

Of course even astrophysics doesn't have a single set of interests or a single set of properties of interest. It may well be that we would need to have multiple modulations depending upon the area of interest (and corresponding microlanguages). This seems to be the conclusion that A'Hearn is ultimately led to:

[I]f ... you want to understand how the interiors of solid bodies work, then you should probably be thinking of Pluto as a planet. If, on the other hand, you want to know how things got to where they are in the solar system, there is no question Pluto got to where it is in exactly the same way as a large fraction of the other trans-Neptunian objects. ... So, if that's the question you're interested in, you absolutely have to classify Pluto as a trans-Neptunian planet. Now, this basically means that you have a dual classification. (Quoted in Tyson 2009: 74)

I want to pause at this point and note that these passages suggest that the modulation should respect the interests of science and scientific properties, and it reasonable to think that scientific efficaciousness is a reasonable criterion for planetariums and scientists to appeal to. But the point needs to be stressed that in other domains—for example,

issues like how to define ‘person’ or ‘rape’—we will need to adjust the definition in response to a different set of needs. That is, the discovery of social and ethical properties will also have to figure in how word meanings are best modulated. The point here is that word meaning modulations need to be responsive to the interests and needs of social institutions (like the scientific community in this case) as our knowledge of the world expands.

(iv) Modulations should not be too taxonomically disruptive.

This was a central argument in avoiding the classification of Pluto as a planet—it simply admitted too many additional objects. Jane Luu, who co-discovered the first Kuiper Belt object, offered this:

We are continuing to try to find more Kuiper belt objects, and the search is going pretty well. What if we find other objects fairly close in size to Pluto—maybe even bigger, or maybe just a bit smaller—will these objects be called planets or what? (Quoted in Tyson 2009: 71)

Or as Michael Brown in the Department of Planetary Sciences at Caltech put it:

Some astronomers have rather desperately attempted to concoct solutions which keep Pluto a planet, but none of these are at all satisfactory, as they also require calling dozens of other objects planets. (Quoted in Weintraub 2007: 227)

Even the *New York Times*, in an editorial published on October 15, 2002, reversed its original position using similar reasoning.

Astronomers predict that they will find up to 10 similar objects in the Kuiper Belt that are as large as or larger than Pluto. So unless we want to add 10 more planets to the elementary-school curriculum, we would be wise to downgrade Pluto to the distant iceball it is.

So far I’ve offered four criteria for reasoning to modulations that I think are reasonable—we could certainly debate their viability further. But the debate about the definition also churned up some proposed criteria which I think are less appealing, or which at least I would want to scrutinize further.

(v) Modulations should allow ease of empirical testing.

This was a criterion that was offered by Weintraub, when he was criticizing definitions based on how an object was formed.

Since we most likely can never know the process by which a free-floating object formed, we would have a very difficult time applying such a criterion to evaluate whether a 10-Jupiter-mass object should be considered a failed star or a large planet. (Weintraub 2007: 211)

His objection is that an object about ten times the size of Jupiter might give rise to deuterium burning and then again it might not, so if we found an object of that size floating free in space we wouldn't know if it was an exhausted star or a large planet. Or probably could not determine from this distance.

I think the criterion is a bit suspect in the first instance because the problematic cases envisioned by Weintraub are far and few between (so far we have no such instances!). So in effect, he is saying that certain definitions should be rejected because in certain rare and so far unencountered cases we would not be able to empirically determine if an object was a star or a planet.

Apart from the rarity of this particular case, is accessibility to empirical test always an important criterion? We can certainly imagine cases where it would be, but if the taxonomy is useful I don't see why we can't live with cases for which identification was in principle not practical. The purpose of the taxonomy is to assist scientific investigation after all, and not to know, for every object we encounter, where it lies in the taxonomy. In some cases it can be just fine not to know, at least for a while.

- (vi) Modulations should not admit relational properties, only individualistic properties.

This is another criterion from Weintraub which I consider suspect. Here Weintraub is taking aim at the third criterion for planet offered by the IAU, which said that a planet had to have swept its orbit clean.

... this now overcomplicated criterion that says that objects that are the largest bodies in unfilled rings are not planets is a flawed means for determining whether an object is a planet, as it does not make reference to the physics of the object itself. (Weintraub 2007: 206)

Whatever we might think about the utility of the third criterion proposed by the IAU, it seems unreasonable that only individualistic properties should be relevant and relational properties not. Science—certainly astrophysics—is neck deep in relational properties, ranging from the notion of an orbit, to gravitational attraction. It seems arbitrary to invoke this criterion without good reason.

(vii) Modulations should be culturally acceptable.

From now on, everyone should ignore the distracting debates of the scientists, and planets in our solar system should be defined not by some attempt at forcing a scientific definition on a thousands-of-years-old cultural term, but by simply embracing culture. Pluto is a planet because culture says it is. (Michael Brown, Planetary Sciences, Caltech, quoted in Weintraub 2007: 226)

This criterion, in my view, is the big loser. If Michael Brown is serious that we should avoid the distracting debates of scientists and return to a thousands-of-years-old cultural term, we would have to go back to calling The Sun and The Moon planets and saying that The Earth is not. We could do that, of course, but I think it is fair to say that there are some unwelcome consequences that would accrue from doing so. Retiring terminology and importing new terminology (like “solar orbitals”) not only comes with some cognitive weight, but it also seems to undermine the role these terms play in current science and science education.

(viii) Modulations should be designed to maximize fun.

In addition, the second [definition] continues to allow the possibility that exploration will find a few more planets, which is a much more exciting prospect than that suggested by the first possibility. We don't think the number of planets found by the current generation of researchers will be large. Maybe one or two more. But we think that letting future generations still have a shot at planet-finding is nice. (Michael Brown, quoted in Weintraub 2007: 227)

Maybe there is something to this criterion that I don't understand, but I have trouble taking it seriously.

The point of this exercise has been to get us thinking about what kinds of criteria are viable in debates about meaning modulation and explicification and what kinds are not. The basic methodology is really

two-layered. First we study the kinds of arguments for word meaning modulation that are on offer, and then we need to reflect on whether those kinds of arguments are normatively viable—that is, whether they serve us well.

Obviously this is just the outline of a research project. One thing that is clear from what we have seen so far is that the normative strategies we use in these cases lean heavily on our abilities at analogical reasoning (Pluto is like the canonical planet in respect Y). Thus we can gain a great deal of insight by studying these debates in the context of work on the psychology of analogical reasoning, particularly in the context of word meaning acquisition—for example, as outlined in Gentner and Rattermann (1991).

I began this exercise with an example ('planet') that is relatively less emotionally charged. What happens when we extend this exercise to more contentious debates on the proper modulation of terms like 'rape' and 'person'?

### 2.3 'Rape'<sup>2</sup>

In the previous section we considered a case in which the meaning of the term 'planet' changed in response to scientific discoveries. While the disputes about the proper modulation deployed a number of arguments, we settled—tentatively—on a handful of criteria that seem normatively apt for these kinds of arguments.

Specifically, we saw that empirical discoveries can serve as triggers for us to re-evaluate our understanding of what a word means, and subsequently to modulate word meanings. When this happens, guiding principles for the modulation include the idea that we should respect core cases, reason analogically from those cases, and track important properties. In the case of 'planet' those basic properties were determined by the interests of astronomical science. In the case of terms like 'rape' the properties in question will be determined by other interests.

<sup>2</sup> See Mason (forthcoming) for more detailed discussion of this topic.

## 52 NORMS OF WORD MEANING LITIGATION

For the discussion that follows I am going to follow the discussion of the evolution of the meaning of ‘rape’ in Schiappa (2003), particularly as it relates to the question of marital rape. As we will see, the case of ‘rape’ tracks that of ‘planet’ in a number of important respects, not least being the idea that meaning modulation should track more fundamental properties and that it should be responsive to relevant empirical discoveries.

As Schiappa notes, initial modulations of ‘rape’ excluded the possibility of marital rape. He notes that in the 1600s Lord Matthew Hale declared that “the husband cannot be guilty of a rape committed by himself upon his lawful wife, for by their mutual matrimonial consent and contract the wife hath given up herself in this kind unto her husband, which she cannot retract” (quoted in Schiappa 2003: 54). As Schiappa (2003: 54) goes on to remark: “Hale has stood as the accepted authority on coerced sex within a marriage” (“To Have” 1986: 1256; Augustine 1991: 560–2). Indeed Hale’s argument is echoed in US judicial decisions well into the 1970s, as in Silberstang (1972: 775):

A husband cannot be guilty of an actual rape, or of an assault with intent to rape his wife even if he has, or attempts to have, sexual intercourse with her forcibly and against her will. The reason for this, it has been said, is that when the woman assumes the marriage relation she gives her consent to marital relations which the law will not permit her to retract in order to charge her husband with the offense.

The mutual consent justification is not the only one that has been given. Also in the mix is the justification given in Sir William Blackstone’s *Commentaries*, published in 1765.

[B]y marriage, the husband and wife are one person in law: that is, the very being or legal existence of the woman is suspended during the marriage, or at least is incorporated and consolidated into that of the husband. (1859: 442)

Both the consent argument and marriage as property argument have echoed in subsequent rulings, including an 1888 ruling by Justice Pollack in *Regina v. Clarence*:

The husband’s connection with his wife is not only lawful, but it is in accordance with the ordinary condition of married life. It is done in pursuance of the

marital contract and of the status which was created by marriage, and the wife as to the connection itself is in a different position from any other woman, for she has not right or power to refuse her consent. (Quoted in “Rape and Battery” 1954: 723n.)

And the argument appears to have held sway until 1977, as evinced in *The State of New Mexico v. Bell*, in which the court argued that the “wife is irrebutably presumed to consent to sexual relations with her husband even if forcible and without consent” (“Rape” 1992: 97).

Whatever trajectory the courts were on in the 1970s, there was a countervailing dialogue taking place outside of the legal realm which pushed back against the legal definitions of ‘rape’, including the seminal work of Brownmiller:

[C]ompulsory sexual intercourse is not a husband’s right in marriage, for such a “right” gives the lie to any concept of equality and human dignity. . . . A sexual assault is an invasion of bodily integrity and a violation of freedom and self-determination wherever it happens to take place, in or out of the marriage bed. (1975: 381)

There are several elements to Brownmiller’s reasoning, but one of her points was clearly that marital rape was *like* recognized cases of rape, in that there was a violation of bodily integrity, a violation of freedom, and a violation of self-determination. Furthermore, marital rape was *like* recognized cases of rape in that it undermines the dignity of the victim of the sexual assault. Of course it is *unlike* other cases of rape in that it happens within the context of marriage, but precisely how important is that fact? The background premise is that the important properties that we want to track in determining the meaning of ‘rape’ should be fundamental social properties like human dignity, freedom, self-determination, bodily integrity—properties which trump the institutional fact that the victim is in a marital relationship with the attacker or that the victim once gave consent.

As in the case of ‘planet’, empirical discoveries have also provided incentive for the broader, more inclusive, modulation of ‘rape’. Part of the motivation for thinking of marital rape differently has been the myth that it is not as damaging or harmful as an attack by a stranger in an alley. But empirical research has blown apart this assumption.

## 54 NORMS OF WORD MEANING LITIGATION

Once the effort was made to listen to victims of marital rape, their accounts revealed that such attacks involved “brutality and terror and violence and humiliation to rival the most graphic stranger rape” (Finkelhor in “To Have” 1986, 1261). Rape by someone supposedly in a loving and caring relationship can be especially devastating. Contrary to the belief that rape victims who know their attackers do not suffer the same sort of ill effects as victims of “stranger rape,” research demonstrates that the short- and long-term effects are typically worse for victims of marital rape (Russell 1982, 190–205; Augustine 1991, 571–72; “To Have” 1986, 1261–62). (Schiappa 2003: 57–8)

In the face of empirical evidence like this, it simply does not make sense to opt for the narrower modulation of ‘rape’; there is good reason to modulate the word to bring more cases within its range. Again, the evidence showed that cases of marital rape were *like* other forms of rape, not just in the loss of freedom and dignity but in the kind of psychological harm done to the victim. That is, they were like recognized cases of rape along a dimension of important properties and facts that motivated our rape laws in the first place.

It is important to understand that, while we are talking about the modulation of a word’s meaning, we are not *merely* talking about word’s meaning; modulations in word meaning have consequences. By modulating the meaning of ‘rape’ to include new cases we bring new individuals under the protection of extant laws. Of course, it could be argued that this is not the best way to go about changing the scope and protection of a law, but this assumes that ossified word meanings are somehow more natural than dynamic word meanings. If I am right, the shifts in word meaning in the legal realm are no different than litigated shifts in meaning in our day-to-day affairs. The original meaning of a word is not privileged, and the decision to privilege it is in fact an active decision to choose a particular modulation of the word while at the same time trying to escape the responsibility of defending the choice of modulation. It is an attempt to assert a position without argument or justification.

Subsequent court decisions (and legislative actions) have slowly come to accept the broader modulation of ‘rape’ to include marital rape, although there has been some recent pushback against this. One of the interesting observations made by Schiappa, is that terms like ‘marital

rape' and 'date rape' have been used like crowbars to help people initially expand the definition of rape. The hope is that eventually, one won't need the prefixes 'marital' and 'date' —these will just be cases of rape.

Prefixes can push things in the opposite direction as well. In recent attempts to carve out exceptions for abortion, some United States congressmen have suggested that abortion should not be permitted in cases of marital rape and date rape. In one sense, this could be taken as just an attempt to carve out exceptions for cases of rape, but many commentators naturally took this to be an attempt to redefine 'rape' —by carving out exceptions we begin to highlight properties that show how these cases of rape are somehow different (they don't involve strangers, for example) and the tacit assumption is of course that these properties are important. In effect, one is pushing for a more narrow modulation of the term 'rape'. The prefixes can be used to broaden meaning but also to narrow it, and we need to be alert to what is going on when they are deployed, since there are likely long-term consequences, for better or for worse.

## 2.4 'Person'

In the case of 'planet' we saw how word meanings can change in response to scientific discoveries, and we have also seen that correct modulation of word meaning should respect those discoveries as well as the properties that are important given the interests and needs of the scientific community. In the case of 'rape' we saw that this general observation can be extended to terms in the social and political realm as well—meanings should be modulated in response to empirical discoveries, and should respect the properties that are important given the interests and needs of our social institutions. In this section I want to examine some relatively more contentious cases involving the terms 'person' and 'human life'. My goal will not be to resolve the issue, but to get clear on the nature of the debate and to outline the form that a productive debate would take.

Once again, I think it is important to understand that when we engage in debates about personhood we are in point of fact engaged

in a debate about the proper modulation of the term ‘person’. The debate is not really about personhood (understood as a debate about something above and beyond the proper modulation of ‘person’), even though philosophy journals are full of articles claiming to probe the concept or nature of personhood.

When we look at the shifting meaning of the term ‘person’ we are going to be interested in how the term has changed in response to empirical discoveries and advances in technology—clearly a big factor in this instance. Like all terms, the meaning of ‘person’ has been and remains underdetermined. What triggered our current debates (what I take to be litigations about the meaning of ‘person’) were technological advances that opened up the door to our having to deal with many murky cases. At the beginning of life we have technological advances that make the survival of a fetus outside the womb more viable, and we also have technological advances that can ensure the health of the mother without the need for abortion. At the end of life, we have technological advances that can keep a person alive after brain death. It’s the usual situation where technological advances and empirical discoveries lead to a rupture in our understanding of what a term should mean. The question is, how can we best litigate the question in this case?

Let’s begin with the discussion of the issue as it was originally framed in the US Supreme Court case *Roe v. Wade* (again I am following the very helpful exposition in Schiappa). Pretty clearly, the question of the range of ‘person’ had significant impact, and this was recognized early on by participants in the court case. The 14th Amendment to the US Constitution says that states may not “deprive any person of life, liberty, property, without due process of law, nor deny any person within its jurisdiction equal protection of the laws.” If a fetus is a person, then it would seem that a fetus deserves equal protection under the law. This was recognized by both sides of the debate. First, it seemed to be conceded by Sarah Weddington, who was arguing for abortion rights, in her exchange with Justice Byron White.

THE COURT: Yes. But I’m just asking you, under the Federal Constitution, is the fetus a person, for the protection of due process?

MRS. WEDDINGTON: All of the cases—the prior history of this statute—the common law history would indicate that it is not. The State has shown no—

THE COURT: Well, what about—would you lose your case if the fetus was a person? ... [I]f it were established that an unborn fetus is a person, with the protection of the Fourteenth Amendment, you would have almost an impossible case here, would you not?

MRS. WEDDINGTON: I would have a very difficult case. (Kurland and Casper 1975: 813–17)

It also seemed to be conceded by the anti-abortion attorney Robert Flowers, in his exchange with Justice White.

THE COURT: Well, if you're correct that the fetus is a person, then I don't suppose you'd have—the State would have great trouble permitting an abortion, would it?

MR. FLOWERS: Yes sir. (Kurland and Casper 1975: 820).

...

THE COURT: The basic constitutional question, initially, is whether or not an unborn fetus is a person, isn't it?

MR. FLOWERS: Yes, sir, and entitled to the constitutional protection. (Kurland and Casper 1975: 827)

THE COURT: Do you think the case is over for you? You've lost your case, then, if the fetus or the embryo is not a person? Is that it?

MR. FLOWERS: Yes sir, I would say so. (Kurland and Casper 1975: 822).

In the Court's ultimate decision, Justice Blackmun drove this point home.

If this suggestion of personhood is established, the appellant's case, of course, collapses, for the fetus' right to life would then be guaranteed specifically by the [Fourteenth] Amendment. (Blackmun 1973: 156–7)

The Court also quickly saw that the Constitution did not provide much in the way of guidance as to what the definition of 'person' should

be—not surprisingly the meaning of the term was underdetermined. As Blackmun put it, “The Constitution does not define ‘person’ in so many words.”

Of course if the question is about whether a fetus falls within the range of ‘person’ the next obvious question is how to go about resolving it. And this is precisely the question that the Court put to the Attorney Robert Flowers.

MR. FLOWERS: [It] is the position of the State of Texas that, upon conception, we have a human being; a person, within the concept of the Constitution of the United States, and that of Texas, also.

THE COURT: Now how should that question be decided? Is it a legal question? A constitutional question? A medical question? A philosophical question? Or, a religious question? Or what is it? (Kurland and Casper 1975: 818)

Indeed, what kind of question is it and how is it to be resolved? I’ve already suggested that it is a *lexical* question, but this doesn’t mitigate the force of the Court’s question, because the correct lexical modulation could depend on whether we are working on a medical context, a religious context, etc. On the other hand, the answer to the question is in a certain sense obvious: We aren’t interested in the proper modulation in all of these contexts—merely in its correct modulation in the *legal* context, which is to say in the context in which we debating whether to extend the range of the predicate to certain individuals in order to bring them under protection of existing laws and constitutionally recognized rights. But how do we answer *that* question?

Justice Blackman (1973: 157) writing in a way that an Original Meaning theorist could appreciate offered the following.

Section 1 of the Fourteenth Amendment contains three references to “person.” The first, in defining “citizens,” speaks of “persons born or naturalized in the United States.” The word also appears both in the Due Process Clause and in the Equal Protection Clause. “Person” is used in other places in the Constitution. . . . But in nearly all these instances, the use of the word is such that it has applicability only postnatally. None indicates, with any assurance, that it has any possible pre-natal application.

All this, together with our observation, *supra*, that throughout the major portion of the 19th century prevailing legal abortion practices were far freer than they are today, persuades us that the word “person,” as used in the Fourteenth Amendment, does not include the unborn. (1973: 157–8)

But here as elsewhere it doesn’t make much sense to divine what the words originally meant—the authors may not have given it any thought, and in any case it is an abrogation of responsibility to fail to ensure the appropriate modulation of critical legal terminology.

Alternatively, the right-to-life lawyers offered a resolution that the meaning of ‘person’ should track important natural properties—the property of being a human being. Thus there was an attempt to anchor the meaning of ‘person’ in a biological category.

Because “person” is not a common medical term, anti-abortion advocates consistently treated certain terms as equivalent: *fetus = live human being = person*. Flower’s statement that, upon conception “we have a human being; a person” indicates that he considers proof of one to be proof of the other. Similarly, the briefs filed by the State of Texas and by various amici curiae (friends of the court) stress such themes as “the human-ness of the fetus,” “the unborn offspring of human parents is an autonomous human being,” and “the unborn person is also a patient.” In these briefs were many photographs of fetuses included to persuade the reader the fetuses, even very early in the gestation period, look like human beings and, thus, should be recognized as persons. (Schiappa 2003: 93)

While it is good practice to try to anchor a definition in more basic and fundamental properties, there is of course the question of whether this happens to be the right set of properties. Certainly, from a biological point of view, the property of being a human being is important, but why should that property carry weight in the realm of law, where we are interested in the plans and goals and interests of agents as they interact with each other and human institutions?

Notice that Schiappa also observes that the briefs attempted to draw analogies between the fetus and uncontroversial cases of persons by showing pictures of the fetus, and demonstrating that they “look like” human beings/persons. Again, whether the argument holds up or not here, this is again a standard strategy in a reasonable definitional dispute. The fetus is *like* a person in that it resembles a person in certain

respects. Of course, it is also *not* like a person in many respects. For example, although the fetus has a “pulse” early on, the heart of the fetus (and its pulse) is unlike an infant in it does not have “the four developed major compartments of the human heart or the developed arteries and veins” (Condit 1990: 212). Likewise, even though an early fetus has measurable “brain waves” and thus is *like* an infant in that respect, the measurable electrical impulses are very *unlike* the brain waves of infants in that not until “somewhere between the twentieth and fortieth weeks do fetuses even begin to have the kind of brain development that would allow perceptions such as awareness of pain” (Condit 1990: 213).

Other legal commentators argued that the fetus as person option blew apart the taxonomy, bringing all sorts of implausible objects into the range of the term (think of the objections to two-part definition of planet based on the fact that many more objects would have to be in the range of ‘planet’).

For example, Chereminsky (1982) observed that, if the Court had held that the fetus is a person, all abortions, even in cases of rape or incest, would have to be prohibited. Indeed, Chereminsky noted that “once it is assumed that the fetus is a person, then there is no legal basis for punishing abortion differently than homicide” (1982: 113). Furthermore, “birth control methods such as the intrauterine device and the ‘morning after pill’ would also be homicide since they act after fertilization and thus kill human lives” (114). Similarly, Tribe suggested that the use of in vitro fertilization would be prohibited since the “process inevitably results in the accidental but foreseeable destruction of at least some of the ova that have been fertilized” and furthermore the Government would be put in the position of regulating pregnancies. For example in the case of a problematical pregnancy the government might have to order the transplantation of the “fetus-person to a less hazardous womb” (1992: 123–5).

Of course, as in the ‘planet’ case, one person’s *modus tollens* is another person’s *modus ponens*, and there are plenty of people in the anti-abortion movement who would sign on to all of these proposals. But just how far would they go? Millions of fertilized eggs are

spontaneously aborted every day. Should we view this as a catastrophic health crisis and immediately invest in research to prevent it from happening?

For that matter, is there any reason that the person should be identified with the fertilized egg as opposed to the egg itself? Mills (2008: 332) has observed that

The sperm and the unfertilized oocyte...are roughly equal insofar as they contribute roughly equally to many salient traits of the later adult. They're not equal, however, in surviving conception. The sperm breaches the egg's cell wall, enters, and *dissolves*. Its dissolution is its death. The sperm doesn't literally exist after conception. The oocyte does. Life is unfair.<sup>3</sup>

One way to think about it is that the sperm is almost like a thumb drive that is used to load a program into the egg and then the thumb drive is destroyed. What about the eggs that don't get fertilized? Each of those is a potential person as well—indeed each of them could have become a person had it been fertilized and allowed to gestate. The *eggs* would have become persons. If this is so then, ought we not to protect all eggs—to make sure that they are all fertilized, then allowed to gestate and be born?

These sorts of considerations have no doubt contributed to courts being shy about basing abortion decisions on the notion of personhood. In stark contrast to the dialogue in *Roe v. Wade*, when the case was reexamined in *Planned Parenthood of Southeastern Pennsylvania v. Robert P. Casey*, then Solicitor General Kenneth Starr backed away from appeals to the notion of personhood. When asked “What is the position of the Department of Justice on the question of whether a fetus is a person within the meaning of the Fourteenth Amendment,” Starr answered that “We do not have a position on that question” (*Official Transcript* 1992: 41–2).

There is much more that could be said on the litigation of definitions still taking place in the abortion debate, but my interest here is not in the rightness or wrongness of abortion so much as in the more narrow question of how the meaning of ‘person’ was litigated within the

<sup>3</sup> Thanks to Keith DeRose for this reference.

## 62 NORMS OF WORD MEANING LITIGATION

abortion debate from *Roe v. Wade* to *Planned Parenthood*. While the issue is obviously as contentious an issue as one can possibly imagine, on the whole the arguments given for one definition over another were reasonable—arguably more reasonable on balance than the disputes in the ‘Pluto’ case. Both sides attempted to argue for the proper modulation of ‘person’ based on arguments from analogy and by appeal to properties that were taken to be important for the case at hand.

So the question is, is there a way to push the discussion forward, or are we better off dropping the question entirely in the case of abortion (as Kenneth Starr apparently did). Starr’s strategy was to shift the debate onto questions that avoid the use of the term ‘person’ as used in the 14th Amendment. That is at best a temporary solution, because whatever may be said about abortion, the definition of ‘person’ is crucial to many applications in the law. This means it is important for the debate (and construction of analogies and disanalogies) to continue, and it is certainly possible if not likely that future empirical (and philosophical) discoveries will inform the question of how best to define ‘person’ in the legal realm.

All this may seem dispiriting to some, because we naturally turn to philosophers to answers to hard problems, but philosophers seldom have the answers. What they do have are principles that should serve us well on the path to finding an answer. My point here is that the crucial principles for resolving the question of what is a person are precisely the principles we ought to use in modulating the meaning of the term ‘person’—or ‘planet’ or any other term, for that matter.

Put another way, our debates about contentious issues such as what is a person are at bottom disputes about how to modulate ‘person’, not about some concept of person in Plato’s heaven. Of course the dispute is not *merely* definitional since there are profound consequences to our choice of modulation and our choice is not arbitrary but founded by important norms for the litigation of word meanings.

This claim may sound tendentious. Is it feasible to think that all moral disputes (or at least many of the key disputes) are metalinguistic? I believe the answer to this is yes. Many (perhaps all) moral disputes are fundamentally metalinguistic disputes (depending on what

we count as a moral dispute; I am not counting disputes that all parties agree turn on empirical facts—for example, on what is the more effective social policy<sup>4</sup>). But the metalinguistic nature of moral disputes doesn't make those disputes less important and (as we have seen in this chapter) it doesn't make them less normative. Much turns on how we define 'person'. Being in the range of 'person' entitles one to a number of rights and protections in both the legal and social realm.

Now it might seem like a debasement of the importance of the debate to say it is metalinguistic, but this reflects confusion about what metalinguistic disputes are—modulation choice is not an arbitrary decision. Everything turns on how we ultimately modulate the meaning of the words. The disputes, although metalinguistic, are no less important. This is a point that has been emphasized in Plunkett and Sundell (2012).

Rather than arguing for the truth or falsity of a literally expressed proposition, speakers engaged in a metalinguistic dispute advocate for their preferred usage of the term given the circumstances. Understanding the meanings of words in terms of the concepts that they express, these disputes involve speakers advocating for using a particular concept by advocating for using the word in question to express that concept. Such negotiations over word usage are largely tacit and center on information that is conveyed via pragmatic rather than semantic mechanisms. And in such disagreements, it is *not* the case that the speakers express the same concepts by the words they utter. Indeed, by definition they do not. We argue that it is possible to understand *many* normative and evaluative disputes as this sort of metalinguistic dispute, while still holding that these disputes express genuine disagreements. Moreover, we argue that metalinguistic disputes can reflect substantive disagreements, disagreements well worth having, and disagreements that would continue to be recognized as such by participants in the dispute even if the metalinguistic nature of their dispute were revealed to them.

As they point out, it is not always transparent to those engaged in a dispute that their dispute is metalinguistic. Indeed, we might add that, far from being obvious to people engaged in meaningful disputes that

<sup>4</sup> Alternatively you might think that in such a case we are morally on the same page but just in disagreement about the best way to achieve some moral end. Imagine two utilitarians arguing about which strategy yields the greatest number of utils. Is this a moral dispute? I would have thought not, but there is no point in making an issue out of this.

## 64 NORMS OF WORD MEANING LITIGATION

those disputes are metalinguistic, the discovery that those disputes *are* metalinguistic is the product of advanced theoretical work in the study of language. How *could* they know?

If this is right, then even apart from cases of moral dispute much of our day-to-day conversation is “meta” —we are not making routine claims about the world but we are making linguistic moves to shift word meanings. It would be an interesting exercise for a linguist to figure out precisely how much of our conversation is meta in this way, but we already know that much of conversation is devoted to conversational repair and explicit requests for word sharpening (specially among philosophers who begin almost every sentence with “what do you mean by . . .”) and in some cases the assertion of semantic authority (as in the Trump dialogue in Chapter 1).

## 2.5 Scalia’s *Original Meaning* Thesis

Presumably, when writing, an author has to anticipate the questions that an interlocutor might ask, and even express those questions in the document (as I have in points throughout this book). Written works also spend more time on the initial explication of terms than would happen in a conversation, where lexical items can be clarified as needed. Still, even after this terminological stage setting, meanings remain underdetermined and there is typically good reason to think that details will have to be fleshed out by readers. In other words, readers can and must continue to modulate word meanings for historical documents.

Over the last few decades, some important legal scholars and judges—most notably US Supreme Court Justice, Antonin Scalia—have made the case that the US Constitution is not a living document, and that we should try to get back to understanding the constitution as it was written by the original framers—sometimes this is called the *doctrine of original meaning*.<sup>5</sup> (Let’s not confuse this with the *original intent* doctrine, which tries to get at what the framers of the constitution *intended* to express.) Scalia’s original meaning theory suggests

<sup>5</sup> Sometimes the doctrine goes by the name ‘textualism’.

that we cannot do better than concentrate on what the constitution actually *says*—on what the words on paper say.

Many of Scalia's formulations of this doctrine come from reports of talks that he has given. So, for example, the right-wing blog *Free Republic* reported on a speech that Scalia gave at Vanderbilt University on April 11, 2005, saying "Words mean what they mean" and "The Constitution is not a living organism." Scalia (1997: 24) offered a more cautious formulation, saying "words do have a limited range of meaning, and no interpretation that goes beyond that range is permissible."

Whatever formulation we go with, pretty clearly Scalia is locked into what I have called the static picture of the lexicon. But 'words mean what they mean' is not the tautology that Scalia seems to think it is. As we have seen, word meanings can change dramatically during the course of a single conversation; how could they not change over the course of centuries? But more tellingly, Scalia's position seems to assume that the original meanings of the words used in the constitution were fully determined—that the meaning of a term like 'person', when used in the constitution was fully fleshed out so that there is a fact about whether it applies to medically viable fetuses, brain-dead humans on life support, and as we perhaps will see in the fullness of time, intelligent robots.

The words used by lawmakers are just as open ended as words used in day-to-day conversation. Indeed many laws are specifically written so as to be open-ended. But even if they were not, there is no way to close the gap and make the meanings of words fully determinate. Technological advances are notorious for exposing the open-endedness of the language in our laws, even when we thought our definitions were airtight. Lawmakers can't anticipate everything. Indeed you could make the case that the whole area of patent law just is the problem of figuring out whether some new object falls within the range of the predicate describing the patented object. Someone makes an object with vacuum tubes calls it a 'blurf', and someone else comes along and makes something very similar with integrated circuits. Does the patent read on this new object? Is it in the range of 'blurf'? Well this is what courts must decide and the idea that the answer is to be found in the language of the patent is, in many cases, absurd.

The problem is that sometimes meanings are not merely underdetermined; they are *wrongly* determined. Our modulations are driven by empirical discovery and sometimes by our better grasp of the important properties underlying the original modulation. For example, we have learned that any reason to take sexual assault by a stranger involving penetration to be ‘rape’ is also a reason to extend the range of ‘rape’ to cover the same acts in a marital context.

Far from being absurd, the idea that the constitution is a “living organism” follows trivially from the fact that the words used in writing the constitution are dynamic and thus “living organisms” in the metaphorical sense in play here. In this respect there is nothing unique about the constitution. It is a dynamic object because of the simple reason that word meanings are dynamic. *Every* written document—indeed every utterance—is a living organism.

I’m not in a position to judge Scalia as a legal scholar, nor do I want to engage his politics here. As the central arguments of this book show, however, his original meaning thesis constitutes a foundation of sand. Furthermore, there is a feature of Scalia’s position which is not merely in error, but which is deeply pernicious. It is deceiving to say that “I am only going by what the document says” when in point of fact there is no stable fact of the matter. I’m all for asserting and defending points of fact, but when one asserts that one is merely going by the letter of a document when there is no static meaning to go by, one is merely taking one’s subjective opinions and wrapping them in the mantles of a sacred document and the fiction that there is a fixed and fully determinate language which settles these matters. In other words, one is supporting one’s position by appeal to an authority that does not exist, in lieu of providing sound arguments and critical thinking.

## 2.6 True When it was Uttered?

There is an interesting question that arises concerning earlier tokenings of words like ‘rape’ under different modulations. For example, what do we say about earlier courts that argued that ‘rape cannot happen in the context of marriage’; were those words true in the mouth

of the judge at the time (because of the meaning the words had at that time) or was the judge saying something false?

This is, I think, a very important question. Suppose a group of individuals forms a microlanguage in which the term ‘athlete’ cannot have women in its range. It sounds horrible to say ‘Women are not athletes’ was true when they said it, or as philosophers like to put it, it was “true in their mouths.”

There are three things to be said about this. The first is that, even if the expression forms were true when uttered, what members of the group said was *wrong* because they had incorrectly modulated the word meaning. (Some philosophers don’t like to think of linguistic forms as the bearers of truth, but we can recast the point this way: The expression expressed a true proposition when uttered. I’ll avoid talk of propositions in what follows, since I don’t think much hangs on which formulation we choose.)

A claim can be wrong because it is literally false but also because it employs an inappropriately modulated term (or both). So, consider the following passage from *Moby Dick* that Chalmers (2011) has highlighted.

I take the good old fashioned ground that the whale is a fish, and call upon holy Jonah to back me. This fundamental thing settled, the next point is, in what internal respect does the whale differ from other fish. Above, Linnaeus has given you those items. But in brief they are these: lungs and warm blood; whereas all other fish are lungless and cold blooded.

Suppose that Ahab and his crew modulate the meaning of ‘fish’ in this way. We might say that ‘A fish is a whale’ is true in their mouths, but what they say is just wrong not because whales don’t fall in the range of ‘fish’ as they have modulated it, but because their modulation is wrong.

We might object to such an utterance by saying that Ahab is wrong or we may object in some other way (we may even say “not true!”). As Plunkett and Sundell (2012) have stressed, even though we are not objecting to the truth of the claim and are in fact objecting to something metalinguistic, our objection is not trivial. When we object because an earlier modulation of ‘rape’ excludes marital rape we are not taking exception to a trivial point of detail; we are objecting to a

modulation that has far-reaching consequences for the welfare of others.

Still, some might object that, even under the cover of calling it a technical philosophical locution, the ‘true in their mouths’ answer doesn’t sit well. To use a vivid example, some Nazis may have modulated the meaning of ‘person’ so as to not include Gypsies and Jews, but we don’t want to be in the position of saying that ‘Gypsies and Jews are not persons’ is false but it was true in the mouths of those Nazis when they said it.

This leads to my second point. There is a reason why the Gypsies and Jews example sounds like it can’t possibly be true in anyone’s mouth—the problem is that single quotes are “leaky.”

Philosophers are trained to use single quotes when they are mentioning an expression, and the conceit is that when something is placed in single quotes we are simply talking about the linguistic form of the expression within it—nothing about the content is of that form is supposed to be relevant. But in practical terms this conceit is misleading. We know, for example, that if an epithet is placed in quotation marks, the quotation marks do not seal off the offensive affect of the epithet. Thus we use expressions like ‘N-word’ rather than the word itself. The offensive affective content leaks out, even when the word is in quotes.

The same is true of semantic content. Brogaard (2008) has observed two respects in which quoted material is leaky. Consider examples (1) and (2).

- (1) ‘I’m going to talk to the doctor’, she said, and she did.
- (2) ‘Give me your money or I’ll shoot’, she said, but I didn’t give it to her.

In the first instance, we have an example of verb phrase (VP) ellipsis, in which ‘she did’ picks up the content of the verb phrase ‘talk to the doctor’. But can that happen if we are to understand the quoted material as merely indicating a string of words? Notice that this example can be made even more effective by embedding an indexical in the quoted material.

- (1+) ‘I’m going to the doctor today’ she said, and she did.

In this instance, I believe the person we are talking about had to go to the doctor on the day she uttered the sentence ‘I’m going to the doctor today’. This suggests that the VP ellipsis is not merely copying the verb phrases from out of the quote, but it is reconstructing the referential content the quoted material had at the time it was uttered.

You might think that we are cheating by using the ‘she said’ locution, but this isn’t necessary to get the example going.

- (1++) She uttered ‘I’m going to the doctor today’ on Wednesday, and true to her word, she did.

I assume this means that she did go to the doctor on the Wednesday she uttered that sentence.

Similarly for (2), it seems that the pronoun is able to pick up the referent of ‘money’; but how is this possible if we are simply referring to quoted material? Notice that this anaphoric leakiness is possible even in the case where we use the ‘true in x’s mouth’ locution:

- (3) In Jane’s mouth ‘Secretariat was an athlete’ is true even though he was a horse.

My interest here is in the ability of the anaphoric pronoun ‘he’ to pick up Secretariat, even though we are supposedly only mentioning the expression ‘Secretariat’ and thus supposedly only concerned with its form. The content still leaks out and is available to the anaphoric pronoun. The upshot is that it is very hard to hear a quoted expression as *merely* quoted.

This leakiness flows in both directions, so that sometimes the content that *we* assign to a term gets attributed to someone about whom we are making a direct discourse report. Consider (4) for example.

- (4) Thales said ‘The planets move’.

It may not be good scholarship to take this as being an attribution of a claim about the planets of our solar system, but it is a very natural attribution for all that. Direct speech attributions can put new contents in the mouths of the person we are attributing the speech to. Now of course philosophers can tighten things up and prevent this from

happening in some restricted contexts, but my point is that it is very natural to slip new contents back into an utterance made earlier, and this explains why it is so offensive to say ‘Jews and Gypsies aren’t persons’ was true in the mouth of a Nazi. We read our own content back into the term ‘person’ even though that term is in single quotes and we aren’t supposed to be talking about its referential content. The term is not only wrongly modulated by the Nazi, but we can’t even bear to use that modulation. The that-clause of our resulting direct discourse attribution is false.

My third point is that this bears certain similarities to the technical notion of relativism about truth as developed in MacFarlane (2003) and others. On that view, truth is relative to a context of assessment, so that it makes no sense to say that the locution was true when it was uttered or true in the mouth of the person who uttered it. All that matters is whether the sentence is true (at that time) given our current context of assessment.

I say that my proposal is similar to this, but it isn’t entirely the same idea. First note that relativism about truth by itself doesn’t solve the problem, as an incorrectly modulated term ‘person’ could steer us to endorse some outrageous claims even in the current context of assessment. So even if one opts for truth-relativism, more needs to be said. If what I am saying is right, not only would the truth evaluation be relative to the context of assessment, but in many cases the meaning of an expression would be relative to a context of assessment.

This latter idea is not particularly new. It comes up in the context of thinking about externalism about content in slow-switching scenarios. In Ludlow (1999) I argued that the contents of our earlier thoughts and utterances may actually shift out from under us as we experience changes in our social environment that impinge upon the meanings of the expressions we use.

As we will see in Chapter 5 I do believe that we have a fair bit of flexibility in discussing alternative modulations of a term. We can put the Nazi’s modulation of the term ‘person’ on the table and discuss it, for example. The problem comes when we start making truth claims about

## NORMS OF WORD MEANING LITIGATION 71

the words uttered by the Nazi, for in those cases we are typically investigating our contents into those words, and when we say that the words are or were true in the mouth of the Nazi, it is hard to avoid hearing it as a claim that it is true in our mouth as well. Flexibility in entertaining meaning modulation does not mean that we can be flexible in the claims and utterances we endorse as true.

# 3

## The Nature of the Dynamic Lexicon

In the previous two chapters we've examined a number of cases of meaning modulation and the elements by which meaning modulation can be normatively guided. We are now in a position to go into more depth on the nature of meaning modulation and the specific theory of modulation I am proposing here.

### 3.1 Some Features of the Dynamic Lexicon

In section 1.1 I argued that we should reject the static picture of language and opt instead for the idea that many of the terms that we use are introduced “on the fly” during individual conversations, and that many familiar expressions have underdetermined meanings that are significantly modulated across conversations and even within conversations.

This dynamic position receives support from work by psycholinguists (e.g. Garrod and Anderson 1987; Brennan 1988; Brennan and Clark 1996; Clark 1992) and their study of lexical “entrainment” — a process whereby the choice and meaning of certain words (sometimes novel words) are worked out on the fly in collaboration with discourse participants.

Psychological studies on entrainment are particularly interesting because they undermine the myth of a common-coin lexicon by showing that even individuals who overhear or witness a conversation are in a much weaker position to understand what is being said than are

the participants. Schober and Clark (1987), for example, show that discourse participants are in a much better position for understanding what is being said because participants are involved in the introduction and modulation of the lexical items that will be employed in the evocation of certain concepts in the conversation.

To see why this should be, think about how much of a lecture you can comprehend by dropping in on a course in the middle of the term. If you are not familiar with the subject matter you may well be quite lost, and not just because you lack familiarity with the objects under discussion (if it is a philosophy class you might have dropped in on an unintelligible discussion of whether tables and chairs exist). One obstacle you may face is that you are unfamiliar with the terminology in play (of course, grasp of the terminology and knowledge of the subject matter are not so easily separated). You were not involved in the process whereby certain terms were introduced into the course. In such situations you may dismiss the terms being used as “jargon,” but this is just a way of saying that you don’t understand the terms being deployed.

My first job after I got my Ph.D. in 1985 was not in academia, but working for the Intelligent Interface Systems Group of the Technology Strategies Center, run by the Honeywell Corporation. My first assignment was to study the then existent machine translation projects—an assignment that sent me traveling to research centers around the world. In those days, machine translation was crude, but in certain circumscribed contexts, it was economically viable to have machines do rough drafts of certain documents. Basically, they did as well as Google translation does today.

Back then, my computer was an Apple II with 48K of ram, and the computers we used at the Technology Strategy Center (Symbolics Lisp Machines) had substantially less power than the low end laptops available for a few hundred dollars today. One might have thought that after twenty years of significant advances in computing power we would also have seen advances in machine translation and natural language “front ends” for databases. But we haven’t. And this is not the least bit surprising.

Most of the work on machine translation and natural language processing has, until recently, been based on a mistake—the idea that one has to find an algorithm that can take some text in a “source language” (for example, English) and in one stroke translate the text into the “target language” (a computer language or another natural language).<sup>1</sup> But this is a confusion from the start.

The next time you go to a bank or a store with a particular request, think about the way your conversation plays out. Do you just make a request and receive an answer? How many times do you have to ask the teller or the clerk to clarify something? (The first time a bank clerk asked “Do you want that large?” I had no idea what she wanted to know.) How many times does the teller or clerk ask you to clarify what you need? How many times do you go back and forth with phrases like “sorry, did you mean . . .” or “I’m sorry, I didn’t catch that” or “I’m not sure what it’s called but I need something that . . .”

There is a great illustration of this from work that was done in the 1980s when we were looking for ways to make computers more user friendly. Before we settled in on the familiar graphical icons, there were attempts to see if text-based commands would work. It turns out that verbal communication with a computer was a difficult problem, but not for the reasons you might suppose.

The real problem was not with the computer, but it was with us and our very shifty and dynamic vocabularies. For example, in studying the way agents attempt to communicate with computers with natural language interfaces, a study by Furnas et al. (1987) found that the likelihood that any two people would produce the same term for the same function ranged from only 7 to 18 percent. For example, when wishing to remove a file, persons used a broad range of terms including *remove*, *delete*, *erase*, *expunge*, *kill*, *omit*, *destroy*, *lose*, *change*, *rid*, and *trash*.

<sup>1</sup> More recently there have been efforts to afford natural language processing systems with this sort of interactive ability. See Purver (2004) for an excellent survey and an introduction to his very interesting program CLARIE, which dialogues with users in order to clarify intended meanings. I believe the strategy remains a minority strategy, however.

You might think you could get around this problem by treating these terms as synonyms and having the system regard any of them as an equally good instruction to delete a file, but Furnas et al. discovered that even with as many as twenty synonyms for a single function, the likelihood of people generating terms from the synonym set for a given function was only about 80 percent. And this is just the beginning of the problem.

When two people *do* use the same term, more likely than not they don't mean the same thing by the term. As Furnas et al. showed, even in a text editor with only twenty-five commands, if two people used the same verbal command, the chances that they intended same function by it was only 15 percent.<sup>2</sup>

In the light of these considerations, think about how silly it was to try and build a machine that “just understands you” when you walk up and begin talking to it. No human can “just understand you” and no machine will ever be able to do it—such a machine is a fantasy machine designed around the myth of a static language. We don't “speak” languages, so if machines did speak static languages that look like English they would be no use in communicating with us anyway. If someone created a static “perfect language” we would have no use for it.

### *Lexical Items Placed in and out of Circulation*

Lexical items are not always in circulation, and indeed, are strategically retired and placed back into circulation depending upon the demands of the microlanguage under construction. The situation is analogous to the position of the traveler who finds that various combinations of US dollars, euros, yen, and Argentinean pesos are accepted in different settings. Some are more widely accepted than others, and some can be introduced in the odd transaction with a bit of cajoling, but at the end of the day there are still establishments where only a peso will do. Lexical items are like this too, but their deployment is more strategic.

<sup>2</sup> See Brennan (1988) for additional discussion of this data.

The experiments on entrainment are particularly illuminating here because they show that additional lexical items are introduced into the microlanguage in response to the need to better discriminate and refine the concepts being deployed. If similar objects are being discussed then there is a greater need to lexically discriminate concepts and kinds of objects and thus there is subsequently increased pressure to introduce more (and more refined) lexical expressions. (This is a point we will return to in sections 5.2 and 5.3 when we discuss the expression of Fregean senses.)

*Meanings of Lexical Items are Underdetermined*

Consider the meaning of the term ‘good’. This is a widely shared lexical item, but there is much to its meaning that is underdetermined. For example, it is a typical phenomenon of sports talk radio to debate which of two sports stars is better. Was Mickey Mantle better than Barry Bonds at baseball? Well, one of them hit more home runs, but the other was on more championship teams. One of them may have cheated by using steroids. Should that be a factor? What is really up for grabs here is the question of what counts as a “good’ baseball player—it is about the meaning of ‘good.’<sup>3</sup>

Jamie Tappenden (1999) offers a formal example of this phenomenon, introducing a language in which some meanings are open-ended and to be sharpened at a later time. The language leaves “certain objects as ‘unsettled’ cases of a given predicate, in that it is open to the speakers of the language to make a further stipulation that the object is, or is not, to be counted as having the property in question.”

As Tappenden notes, these cases happen frequently both unintentionally and intentionally outside of formal languages, with an example of intentional cases coming from the realm of law:

This happens with some frequency in law: it may be convenient to stipulate a condition for only a restricted range, leaving further stipulation for the

<sup>3</sup> Possibly debates like this that involve ranking individuals are a way of establishing a scalar meaning of ‘good’ as applied to some relevant class—in this case baseball players. For a very interesting discussion on how word meanings typically incorporate context-sensitive scales (meaning the ordering on the scale changes according to context), see Barsalou (1987).

future. There have been many different reasons for such reticence: courts have wanted to see how partial decisions fly before resolving further cases, higher courts may want to allow lower courts flexibility in addressing unexpected situations, legislatures may be unable to come to the needed political compromises without leaving ‘blanks’ for courts to fill in.<sup>4</sup>

Tappenden is thinking of cases in which matters are intentionally left open, but we can imagine lots of reasons why aspects of word meaning might remain open as a kind of natural default state—it may simply be too costly to determine everything (even for an expert) or it may be that crucial aspects of word meaning depend upon the discourse situation and/or facts about the world that remain open.

Another good example of this that we discussed in the previous chapter is the recent debate about whether ‘planet’ should be modulated so as to include Pluto. In 2003, the IAU Working Group on Extrasolar Planets put the situation this way.<sup>5</sup>

Rather than try to construct a detailed definition of a planet which is designed to cover all future possibilities, the WGESp has agreed to restrict itself to developing a working definition applicable to the cases where there already are claimed detections. . . . As new claims are made in the future, the WGESp will weigh their individual merits and circumstances, and will try to fit the new objects into the WGESp definition of a “planet,” revising this definition as necessary. This is a gradualist approach with an evolving definition, guided by the observations that will decide all in the end.

As should be clear from earlier chapters in this book, I believe these situations to be ubiquitous. It is not merely when we are engaged in law and astronomy that we use terms that have underdetermined meanings; we use such terms all the time.

Typically, it doesn’t matter that our word meanings are underdetermined because we are using expressions in circumstances in which it is clear how to determine whether the predicate applies to an individual or not. If it *isn’t* clear whether it should apply, we modulate the meaning until it is clear whether the predicate applies or not.

<sup>4</sup> See also Endicott (2000) on this point.

<sup>5</sup> The passage specifically talks about definitions or, if you will, explicifications, but clearly underlying this is the question of how the word should be modulated.

I believe that it would be a mistake to try and assimilate these cases of underdetermined meanings to those of vague predicates like ‘bald’. Many of the disputes that arise have little to do with vagueness. Consider the dispute I heard on WFAN (a sports talk radio station in New York) when *Sports Illustrated* announced its “50 greatest athletes of the 20th Century.” Some listeners called in complaining that a horse—Secretariat—had made the list, while host Chris Russo defended the choice. Clearly this is a dispute about what should be in the extension of ‘athlete’, and the callers wanted to argue that a horse had no place here. It is not as though the dispute would be resolved if Secretariat were a little bit faster or could throw a baseball, so it seems hard to imagine that these are vagueness cases.<sup>6</sup>

This is also a good example of a case where fleshing out the meaning of the term is up to us and our communicative partners. So, even when we are deploying a common lexical item (like ‘athlete’, for example) the range of the term within a given context may be up for grabs and may require some form of coordination strategy—in the sports talk radio case the coordination took the form of a debate where discourse participants argued their respective cases.

At least in this narrow instance there is an obvious similarity to the legal realm, where competing parties may come together to resolve a dispute—in this case the way in which the term is to be understood with respect to the new cases in question; think of question of whether an existing patent “reads on” (applies to) some new technology. The key difference is that rather than taking place in a formal courtroom setting, these debates play out in less formal realms, ranging from sports talk radio to arguments with colleagues, friends, and partners.<sup>7</sup>

<sup>6</sup> John Hawthorne suggested to me that maybe the relevant Sorites scale is along another dimension: “maybe if Secretariat were a bit more human.” Suppose Secretariat could talk like Mr Ed, do basic math, run on two legs, etc. Would this change minds? I think not.

<sup>7</sup> There are technical issues that I am avoiding here, not least of which is the logic of underspecification. How do inferences work when meanings are incomplete or underspecified? For a sample of work on this topic see van Deemter and Peters (1997). I will examine this issue in more detail in Chapter 4.

*Assigning Meanings to Lexical Items by Jurisdiction*<sup>8</sup>

Tappenden's metaphor of court decisions can be extended in fruitful ways. Disputes over the best baseball player or whether a horse counts as an athlete are often just wheel spinning, but sometimes a consensus is achieved. This might be due to a series of rational arguments or it might be a simple case of someone asserting a claim and other participants deferring. It is worth noting that when these disputes are resolved there are often jurisdictional limits.

When courts come to a decision on a particular dispute they set a precedent that may carry over into other jurisdictions. On the other hand it may not. Similarly, we may resolve a dispute or coordinate on the meaning of a term, and expect that to carry over into other micro-languages that we form. We may be disappointed to find we have to reargue our point of view, or reestablish our credentials.

Alternatively, it may be that some of the disputes that we engage in (about sports, television, movies, and questions like "Is Chesner a smoker if he only smokes when drinking?") which *appear* trivial or silly are valuable precisely because they are litigating the content of certain key terms and this may be valuable in contexts where more is at stake and communication is critical. In other words, idle talk may well serve the function of helping us to calibrate our lexicons during periods of down time. These periods of calibration may serve us well later when we need to collaborate on some important project or problem.

Sometimes we may not be involved in litigating the meaning of a term, but we may rather defer to someone else's usage (perhaps in the conversation, or perhaps in the greater community). To use a famous example from Putnam, we may defer to an expert on the proper individuating conditions of the expressions 'beech tree' and 'elm tree'. There may be a social division of labor involved in fixing the semantic content of our utterances.

<sup>8</sup> This section was inspired by discussions with Henry Jackman.

### 3.2 Core Ideas

I've canvassed a number of examples of meaning modulation, provided a picture of how meaning modulation can be normatively guiding, and in the previous section I gave a general gloss on how meaning underdetermination and meaning modulation work. It is probably a good time to pause and lay out the central doctrines that are in play, and there are quite a few.

#### *Meaning Underdetermination*

In Chapter 1 I gave the example of someone asking me how many books I had written, and me responding by asking what they meant by 'book'. I had to ask the question because the meaning of 'book' is *underdetermined*. Nothing in our broadly shared understanding of the meaning of 'book' was sufficient for me to answer the question. Thus we had to modulate the meaning to the point where I could answer the question. As I also noted in the introduction, I think this is a pervasive fact about the words we use: They are all underdetermined to some extent or another.

#### *Representational Austerity and Representational Neutrality*

We need to be careful to distinguish between word meanings and the way that those word meanings are represented. It is entirely possible (probable, I should think) that the meaning of a word is not exhausted by the way in which we represent it. How much is represented? By some accounts, not very much. A good case can be made that the part of the meaning explicitly represented in lexical items consists of just hints and clues (like one finds in dictionary entries) that may help us to deploy cognitive resources (typically analogical reasoning) to flesh out the word meanings, and the way we flesh them out will vary according to contexts and social settings. A similar point has been made by Rayo (2013), who talks about a "grab bag" theory in which we rely on a variety of resources to work out word meanings. As Putnam has stressed, elements of meaning could depend upon other members of our linguistic community and likewise facts about the world. Let's call the view that representations of meaning underspecify meanings *Representational Austerity*.

As intuitive as Representational Austerity may seem, I am not climbing on board with it here because I am suspicious of there being a clear idea of what representations look like and how austere or robust they are. One idea would be that a representation is a data structure—a semi-stable, localizable structure in the computational state realized by an agent (person or machine, for example). But as I argued in Ludlow (2011: ch. 2) I seriously doubt that there is a brute fact about systems individuated narrowly that determines what computational states they realize. Furthermore, whatever states they realize and however those states are individuated, the connection between those states and what they represent is largely mysterious. There is no reason to think that the data structure has to be isomorphic to what is represented in order for us to represent something (a meaning, for example). So what then is the connection? There is hopefully an answer to the question, but the doctrine of *Representational Neutrality* says that the theory of the dynamic lexicon is neutral on the matter.

Accordingly, in what follows, I'm not going to have much to say about how meanings are represented. Instead, I will talk directly about meanings, whether they can be broadened and narrowed, whether they are fully determinate, whether we can control them, and ways in which we make them explicit without talking about the role of representation in this. This is not to say that meanings aren't represented or that the representation of meaning is not part of a full story of meaning; it is just to say that the theory I am proposing here is neutral on all of these matters.

This claim of neutrality may seem anomalous when we begin talking about, for example, the thematic structure of words (e.g. meaning facts involving roles like agent, patient, theme), but the anomaly is only apparent. One can perfectly well subscribe to these thematic relations without taking a stand on how these roles are to be represented (or even on whether they are represented at all).

### *Meaning Egalitarianism*

Sometimes people argue that there is a primary or privileged meaning for a term. It is difficult to see what this might be in the case of 'book'

but in other cases the intent is clear enough. The privileged/primary meaning of ‘flat’ would be a meaning like absolutely flat. The privileged/primary meaning of ‘know’ would be something like knowing without any possibility of error.

Sometimes people argue that expressions like ‘six foot six’ or ‘3 o’clock’ have privileged meanings that are their “on the nose” meanings. When we say ‘we will be there at 3:00’, there is a privileged/primary meaning of that expression which means precisely 3 o’clock down to the nanosecond. We may be entitled to diverge from that meaning, but the core meaning is the absolute or on the nose meaning, and any divergence from that core meaning is just that—divergent.

*Meaning Egalitarianism* is the doctrine that there are no privileged/primary meanings—no “absolute” senses like absolutely flat or “on the nose” senses like precisely six foot six. The absolute and on the nose senses are simply two modulations among many others that are, at the outset, equals. The absolute and on-the-nose senses do not come first, and they need not be the goal of a proper lexical modulation. This does not mean that all modulations are equally good. As we saw in Chapter 2, there are norms governing how we modulate word meanings, and the best modulations will turn on our interests and the important properties that anchor our interests. They will not turn on the approximation of some absolute or on-the-nose meaning. In fact, such meanings may not even exist, which leads to the next doctrine.

### *Meaning Imperfection*

As I just noted, some people believe that the core meaning of a term like ‘flat’ is a sense like absolutely flat and that the core meaning of an expression like ‘six foot six’ or ‘3:00 on the nose’ are precise spatial lengths and temporal locations. As I said, meaning egalitarianism rejects the idea that these are the core meanings, but the doctrine of *meaning imperfection* suggests that such meanings may not exist at all. Is there a coherent notion of absolute flatness or exactly six foot six in mathematics or physics? There is certainly reason to doubt it.

Just to illustrate, consider a predicate like ‘flat’. Obviously there are no surfaces that are absolutely flat in the physical world (presumably

everything is bumpy at the micro level), but it isn't clear that there is a stable notion of flatness in the mathematical realm. We might think, for example, that a two-dimensional plane is absolutely flat, but what then happens when the plane is in a non-Euclidean geometry. In some such geometries a plane can be quite "bumpy." Of course we could stipulate that a plane in a particular axiomatic system like Euclid's is flat, but this misses the point. The question is whether there is a notion of absolute flatness (or absolute anything) that is stable across contexts. The doctrine of meaning imperfection says that there isn't.

### *Meaning Control*

The doctrine of *Meaning Control* says that we (and our conversational partners) in principle have control over what our words mean. The meanings are not fixed by convention, nor by our conversational setting alone. If our conversational partners are willing to go with us, we can modulate word meanings as we see fit. This isn't the Humpty Dumpty theory of word meaning because Humpty needs Alice to play along. If she doesn't play along, the word meanings are not modulated as Humpty wishes. It does not follow, however, that there isn't a right way and a wrong way to modulate word meanings. As we saw in Chapter 2, there are still norms that govern this process of modulation.

Meaning Control does not exclude the possibility of externalism about content—either environmental externalism or social externalism as in Burge (1979)—nor does it preclude the possibility of a division of linguistic labor as in Putnam (1975). The idea is that it is within our control to defer to others on elements of the meaning of our words (for example, a doctor on the meaning of 'arthritis' and a botanist on the referents of 'beech' and 'elm') and it is also within our control to be receptive to discoveries about the underlying physical structure of the things we refer to (for example, the discovery that 'water' refers to  $H_2O$  and not XYZ). The theory of the dynamic lexicon is largely neutral on theories of content; my point here is simply that the dynamic lexicon is compatible with existing theories that employ some version of

externalism about content. And conveniently so, since I happen to be an externalist about content.

### *Concepts as Linguistic Meanings*

Clark (1992) talks about the result of lexical entrainment as being a “conceptual pact.” I rather like this idea. The basic thrust of it is that concepts are not things floating in Plato’s heaven or etched in some data structure of our mind/brains; rather they are simply word meanings. That by itself may sound like a standard thing to say (one could reverse the equivalence and argue that word meanings are objects in Plato’s heaven or are things etched in data structures in the mind/brain after all), but when it is linked with the idea that we collaborate to modulate word meanings it suggests that concepts are the sorts of things that we build together on the fly. I won’t be defending this thesis in the book, but it is a doctrine worth some reflection along the way: Concepts are underdetermined, modulable, and often the product of collaborative effort.

## 3.3 Terminology

I’ve already alluded to some of the following distinctions, but it will be useful to develop them in a bit more detail and set them out in one place.

### *1 Vagueness vs. Ambiguity*

I understand vagueness and ambiguity to be *prima facie* distinct phenomena. Let’s focus on the word form ‘bank’ and consider the meaning associated with financial institutions and the meaning associated with river embankments. The distinction between those two meanings is an instance of ambiguity. Vagueness has more to do with uncertain boundaries. This is not to say that ‘bank’ isn’t vague. Once disambiguated, the two separated meanings remain vague (there are borderline cases of financial institution banks and borderline cases of river banks). So, ‘bank’ is both ambiguous and (once disambiguated) vague.

## 2 *Vagueness vs. Variable Comparison Classes*

The adjective ‘large’ is sensitive to comparison classes (large for an elephant vs. large for a flea) but this is not vagueness, as I understand it. Of course, ‘large’ is vague—once we fix the comparison class to elephants we still have unclear borders between the large elephants and not large elephants.

## 3 *Meaning Underdetermination vs. Vagueness*

As mentioned earlier in this chapter, several years ago *Sports Illustrated* produced a list of the top athletes of the 20th century, to which they added the famous racehorse Secretariat. I was listening to a sports talk radio show when viewers called in to complain that horses can’t be athletes. Arguments ensued, but I didn’t take them to be arguments over a borderline case. Had Secretariat been a centaur (and hence arguably part human) we might have had such an argument. These were rather arguments about whether ‘athlete’ should be defined in such a way as to admit horses.

We can, however, say that vagueness is a *kind* of underdetermination—it is underdetermination along a scalar dimension.<sup>9</sup> So, for example, we can argue over whether I am bald, but in that instance the argument seems to be about whether I am far enough along on the baldness scale to count as bald (here assuming the border is underdetermined).

## 4 *Meaning Underdetermination vs. Meaning Indeterminacy*

Meaning Indeterminacy is of course the notion that Quine was introducing with his famous ‘gavagai’ problem. We parachute into an exotic culture and then a native says ‘gavagai’ when a rabbit hops by. Did the native mean rabbit by this expression, or did she perhaps mean rabbit stage, or rabbit parts, etc.

Meaning underdetermination is different from this because in this instance the behavior of the native is consistent with rabbit (or stage or

<sup>9</sup> I owe this observation to Rebecca Mason.

parts thereof) being a clear instance of ‘gavagai’. Whatever she was gesturing at was clearly (to the native) part of the range of the expression. Cases of underdetermination arise when we point to something and ask, in whatever the native question intonation is, “gavagai?” and our interlocutor is stumped or the natives begin arguing about the matter before resolving it. I’m not taking a stance here on whether meaning indeterminacy is a serious philosophical problem; I’m merely saying it isn’t the same thing as meaning underdetermination.

*5 Meaning Underdetermination vs. Meaning Underspecification*

Meaning underspecification typically involves the idea that it is not specified which of several determinate meanings are intended. For example, one way of treating ambiguity would be by saying that a term like ‘bank’ has a single lexical entry with an underspecified meaning because it hasn’t been *specified* which of the typical meanings is intended. That is, there are two quite clear meanings for ‘bank’, and the lexical entry for ‘bank’ is indifferent between them. Thus underspecification suggests that there is a fully fleshed out candidate meaning out there, but that the lexical entry fails to pick one. The thesis of meaning underdetermination is different from this because it does not suppose that there are determinate word meanings out there for us to (un)successfully specify—word meanings are, as it were, yet to be fleshed out and never will be fully fleshed out.

Now clearly this point is controversial. One might, for example, argue that there are many possible languages (understood as abstracta) and it is underspecified which of them we are speaking. Meaning modulation would then be a process by which we change the set of possible languages. Narrowing a meaning would be to narrow the class of languages.<sup>10</sup>

I don’t think this strategy works. That is, I don’t think that there is a precise set of languages that represents the possible meanings of

<sup>10</sup> See Lewis (1972). Thanks to Ernie Lepore for discussion here.

an expression. It will take some work to show this, however. For now I simply want to point out that, with provisos, underdetermination and underspecification are different things.

### *6 Open Contextual Parameters vs. Meaning Underdetermination*

Some people think that verbs (and perhaps also nouns) have open variable positions or parameters, to be filled by time, location, agent, etc. So, consider an expression like ‘it’s raining’. On such a view, an utterance of this expression has open parameters for place and time. Whether or not it is a correct analysis, I consider these open contextual parameters to be distinct from meaning underdetermination. It is also often argued that indexical expressions like ‘I’ and ‘now’ have open contextual parameters—the meaning of an utterance of ‘I’ depends on who is doing the uttering. Again, this does not count as meaning underdetermination.

On the other hand, this is not to say that expressions like ‘I’ are not underdetermined. Note that there is still underdetermination once the parameter is fixed. That is, if the speaker who utters ‘I’ is me, the parameter is set to me, but does that include my hair? Clothes? Glasses? Hands? In the game of baseball a ball can “hit me” by hitting only my shirt or hitting my hands, but *I* can still lose both my shirt and hands and still be me. Similarly, if I utter ‘now’ the time interval is fixed to the time of utterance, but is that this minute? This nanosecond? Today? This geological era? Again, I take this to be underdetermined.

### *7 Sharpening Meaning vs. Narrowing Meaning*

Modulation is the mechanism by which meanings can be narrowed and broadened on a conversation-by-conversation basis. This is not the same thing as saying that the meanings are being sharpened and unsharpened. The reason is that we can narrow the range of things to which a term applies without sharpening the term; that is, we can narrow the range of things to which a term applies without making the borderline cases sharp. To follow up on our example, we can agree that horses cannot be athletes without thereby sharpening the definition of

‘athlete’ —there are still just as many borderline cases once horses are ruled out. The same holds if we broaden the definition of ‘athlete’ — for example to include grandmaster-level chess players. This doesn’t make the definition of athlete less precise. It just means that the range of things to which the predicate applies has been broadened.

### 8 *Being a Semantic Value of a Predicate vs. Being in the Extension of a Predicate*

Elsewhere in this book I have been careful to avoid talking about the *extension* of predicates like ‘athlete’, and have spoken about being in the range of the predicate instead. I don’t intend ‘range’ to be another way of saying ‘extension’. An individual is a semantic value of a predicate just in case the predicate is true of that individual. As we will see, what counts as a semantic value of a predicate (and what a predicate is true of) is underdetermined. I take an extension to be a fixed set of entities, but for reasons we will see later, I don’t think there is a determinate set of objects that are semantic values of natural language predicates.<sup>11</sup>

### 9 *Explicifying vs. Sharpening*

‘Explicifying’ is my word for introducing an explicit definitional component to a word meaning. For example, we might stipulate that cars can’t be athletes. That doesn’t really narrow the meaning because no one actually supposed they were, nor does it sharpen the meaning because it doesn’t really help us with any of our borderline cases. Still, it adds an explicit definitional component to the meaning that we may or may not choose to take on. Obviously explicifying happens more frequently in institutional settings. It also happens in the beginning of academic publications—as in this section of this book.

I don’t mean to suggest that an explicification must be written or externalized. It could be common knowledge to discourse participants

<sup>11</sup> I am indebted to Chris Gauker for discussion here. For a similar point, see Gauker (2002). It is of course possible to modulate the meaning of ‘extension’ so that extensions (and I suppose therefore sets) are underdetermined but doing so might be a recipe for confusion.

without being voiced. I assume that if one has an explicification one can express it linguistically if called upon to do so.

### 10 *Sharpening*

Given what I have said about sharpening, it should be clear that I intend it to be a very context-sensitive notion. To sharpen a meaning is to modulate it in a way that avoids borderline cases in that context. Thus a meaning that is sharp in one context with a few borderline cases could fail to be sharp in a context with lots of borderline cases.

Suppose we took a term like ‘tall’ and specified that it meant over six feet tall. Isn’t this a sharpening of the definition of ‘tall’? On my view it is certainly an explicification of ‘tall’ but whether it is also a sharpening depends on whether it helps us avoid borderline cases. In most cases it would count as a sharpening, but in a context where many people have a height near six feet tall within the range of error of our measuring device it is not a sharpening.

The case of ‘planet’ discussed in Chapter 2 provides a good example of the context sensitivity of ‘sharpening’. The third clause for the definition of ‘planet’ says that to be a planet you have to have swept your orbit of debris. That definition is sharp enough in our solar system, where the first eight planets have relatively clean orbits (or at least they do now), but what about a solar system in which the orbits of the planets have significantly more debris, so that each planet is a borderline case. In such a solar system the three-clause definition of planet is not sharp.

### 11 *Definitions*

I haven’t said much about definitions, but for the record I take a definition to be the result of the explicification of the meaning of a term. Typically, it is an agreed-to explicification or is at least proffered as something that can be agreed to.

### 12 *‘Meaning’*

I haven’t offered a definition (explicification) of ‘meaning’ and I don’t intend to. Suffice it to say that there are numerous theories of what ‘meaning’ means. My view is that the meaning of ‘meaning’

is underdetermined and the goal of this book is not to sharpen that meaning but rather to offer a theory of how word meanings are modulated while at the same time being as neutral as possible on what one's theory of meaning should look like. My hope is that the proposals made here will have uptake in a broad range of theoretical frameworks. For example, I would like this proposal to largely be neutral between approaches that take meaning to be largely semantic and those that take it to be largely pragmatic. Likewise it is neutral on how meanings are represented and even between accounts that traffic in wide content and those that utilize narrow content. The key point is that the account of modulation and meaning underdetermination will look very much the same from all these perspectives.

### 3.4 Precursors

The theory I develop here shares common ground with ideas that have been put forward by 20th-century and contemporary philosophers and linguists, and it will be useful to go over the common ground before I start drawing distinctions and separating my view from these other projects. To keep the discussion clear, I begin with the traditional distinction between phonology, syntax, semantics, and pragmatics. To a first approximation, we can say that phonology is concerned with the representation of the features that impinge on how we perceive and articulate linguistic forms. Syntax has to do with the structural form of language (for example, how a well-formed sentence is composed of grammatical categories like noun phrase and verb phrase), semantics with the interpretation of those forms, and pragmatics with how those interpreted forms are used and understood in particular contexts.

Not everyone has seen a place for all of these components of grammar. Beginning with the later Wittgenstein (1958), philosophers have stressed the importance of the pragmatic component of language in understanding linguistic utterances. In Wittgenstein's case, meaning is tied to use which in turn is tied to a "form of life," which crudely can be understood as a context of language use. It is reasonable to think that the later Wittgenstein saw no place for a semantic component to

the grammar, and it is arguable that he likewise had doubts about the coherence of the syntactic component.<sup>12</sup> We can call this view “hyper-contextualism” (a more recent representative of this view would be Travis 1985, 1989, 1996).

One way of thinking about the hypercontextualist approach is that there is phonology and there is pragmatics (in this case a use theory), and not much in between (syntax and semantics drop out of the picture). Others have argued that there is phonology, syntax, and pragmatics, but not much role for semantics.

### *The Role of Semantics*

Grice (1989) pushed back against this kind of hypercontextualism with an extremely elegant theory that showed how one could have a well-defined semantic component, and a reasonably well-understood pragmatic component, and use the two in conjunction to construct a viable theory of meaning (i.e. theory of “what is meant”).<sup>13</sup>

Grice’s famous example of how this worked involved an instance of damning with faint praise. Suppose that I have a student—let’s call him Bob—who is not especially bright but who comes to me asking for a letter of recommendation. I don’t wish to speak negatively of Bob, but on the other hand I don’t want to write a dishonestly positive letter for him. So in my letter I simply write, “Bob is always very punctual and has very good handwriting.” The interesting thing to notice here is that what I literally said was something about his punctuality and handwriting, but what I was trying to communicate was something more—I was communicating the thought that this would not be an appropriate hire.

Grice held that the move from what is literally said to what is meant involved the deployment of rational communicative strategies—strategies that were encoded as his “cooperative principle” in the form of maxims: the maxim of quantity (say as much as needed and no more),

<sup>12</sup> This would be the case e.g. on Kripke’s (1982) interpretation of Wittgenstein’s (1956, 1958) rule-following argument.

<sup>13</sup> See Neale (1992).

maxim of quality (try to make your contribution one that is true), maxim of relation (be relevant), and the maxim of manner (be perspicuous). One could take what was literally said, apply the cooperative principle and maxims, and work out what was meant. In the case of the handwriting example, I was flouting the maxim of relevance and the maxim of quantity, so I must have been trying to communicate that I lacked something positive to say about Bob, or perhaps I had something negative to say that I did not wish to commit to paper.

Another good example of the pushback against hypercontextualism involves a response to Wittgenstein's thesis about knowledge claims in *On Certainty* (1969). Wittgenstein observed that when we make knowledge claims like 'I am certain' or 'I know' it typically involves a context in which there is some doubt about the claim, so we shouldn't think that certainty involves the usual philosophical analysis of knowledge as justified true belief. But as King and Stanley (2005) have observed there is a natural Gricean response to the Wittgensteinian analysis—we typically don't make knowledge claims when there is no doubt about the matter because doing so would flout the maxim of quantity—don't say more than is needed given the conversational goals.

Much recent work in semantics and pragmatics has not disputed that there is a distinctive role for both semantics and pragmatics, but has taken issue with Grice on the particulars. For example, one interesting development in the post-Gricean pragmatics literature has gone under the banner "Relevance Theory" (Sperber and Wilson 1986; Carston 1997, 2002) and it has employed several key ideas that are worth our consideration. One idea is the thought that the driving principle that gets us from what is literally said to what is meant is not a theory of rational communicative principles (as Grice thought), but rather simply relevance—which is not Grice's notion of relevance but is at bottom a kind of "least effort" principle.<sup>14</sup> Communication requires the expenditure of energy and so does comprehension. Efficient communication requires that we package what we say so that

<sup>14</sup> Thanks to Dan Sperber for very helpful discussion here.

our communicative partners can unpack what we say with minimal effort. It is a kind of optimizing problem that seeks to minimize the effort expended by the encoder and decoder of the message.

To illustrate, the relevance theorist might say that the reason we don't make knowledge claims in the case where we are certain of some claim is simply that it is a waste of energy to do so. Communication is expensive. Thus the maxim of quantity is subsumed under the principle of relevance—again understood as a least effort principle.

Relevance Theory thus contrasts with Grice's approach in two ways. First, it subsumes all of Grice's maxims under the "least effort" principle (which they confusingly call "relevance"), and second it represents a shift away from thinking about pragmatics as a process involving rational communicative principles and repositions it as a lower level process in cognitive psychology. For Grice, the maxims were normative. In Relevance Theory, the principle of relevance is only normative in the sense a biological norm is. Relevance Theory is, at bottom, a descriptive project and not a normative project.<sup>15</sup>

But there is another key departure from Grice in Relevance Theory. For Grice, the semantic component—which delivered what is literally said—was more or less independent of pragmatics. Of course, one would have to have some contextual inputs to the semantics (the reference of indexical expressions like 'I' and 'here', for example, are determined by who is uttering them and where) but once those pre-semantic components are fixed, the semantics could proceed without fear of "pragmatic intrusion" (an expression from King and Stanley 2005). Relevance theorists and others have stressed that the contribution of the semantics is more impoverished than Grice seems to have supposed—we often make utterances that are incomplete propositions, and often use expressions with meanings that are only partially encoded. We then rely on a process of "free enrichment" to flesh out the proposition meant.

<sup>15</sup> Thanks to Dan Sperber and Deirdre Wilson for discussion of this point.

Relevance theorists put the project this way. Word meanings are only partially encoded in our speech (fully encoding our message is costly). Pragmatics, in the form of Relevance Theory takes the information that is explicitly encoded and utilizes two processes to work out the intended meaning of the utterance. One process is *explicature*—spelling out the full logical form and literal meaning of the expression (for example by filling in ellipsed material). The other process is *inferential*—deducing what is meant. These two processes work in parallel. We make certain inferences about what is intended (based on least effort assumptions) and these help us to flesh out the partially encoded message. This in turn allows us to make more inferences about the intended meaning.

The Relevance Theory choice of the term ‘encoded’ is unfortunate I think, because in communication theory the information encoded precisely is the information that can be extracted from the signal, no matter how impoverished the signal. I think a better way of making the point that relevance theorists want to make is the following: The information provided by the semantics is *partial* and often *sub-propositional*. To use an example from Stainton (2005), I might pick up a letter and say “from Paris.” Most relevance theorists would suggest that what the semantics delivers here is quite thin. The intended message must be extracted from this thin semantic contribution utilizing the processes of explicature and inference. Others, like Recanati (1992, 2004, 2010) have argued that the semantic contribution is virtually negligible—it is pragmatics all the way down, although Recanati holds that the process by which we pragmatically construct the content is amenable to theory construction.

The approach I am taking in this book is not a *radical* departure from this general picture, but it is a departure on a number of important points of detail. *Prima facie*, I see no reason why the moves I am making cannot be incorporated into the general relevance theory project. Of course, both God and the devil are in the details.

Here is one such detail: I believe that most relevance theorists believe that the thought being expressed is fully determinate. It is my view that if there is no interesting difference between the vehicles of

thought and the vehicles we use to express thought, I expect thoughts to be underdetermined in all the ways that language is, as outlined earlier in this chapter.<sup>16</sup>

Here is another detail: I believe that there is more to syntax and semantics than meets the eye, and I actually believe that an utterance like ‘from Paris’ might well have a full (if unpronounced) clausal structure. I don’t think it is implausible to suppose that the semantics can deliver a proposition from such an utterance (see Ludlow 2005a, 2006a). This, however, is a complicated technical matter of bookkeeping and it needn’t distract us here. Most parties to the dispute (perhaps all) agree that this is an empirical question and not a conceptual one.

I am inclined to agree with relevance theorists that there is a great deal of contextual effect on how a sentence utterance is processed. However, the fact that contextual factors influence how we interpret or parse a sentence seems to me completely obvious and, for that matter, benign. I believe that pragmatic processes even influence our auditory perception of phonological features, word boundaries, and, below that, morphological bracketing (contrast ‘the door is unlockable if you have the right key’ and ‘there is no way to secure this room, the door is unlockable’ —the contrast is between [[unlock]able] and [un[lockable]]).

The point here is that I believe that the intrusion of real world facts into phonological, syntactic, and semantic perception do not undercut the idea that phonology, syntax, and semantics are separate modules in our cognitive psychology. Being subject to pragmatic intrusion does not make something part of pragmatics.

In this vein, King and Stanley (2007: 230–1) offer a helpful way of illustrating the point.

we can distinguish between two ways in which context determines what is communicated. The first way context may determine what is communicated is by affecting the semantic content, via resolution of the referential content of context-sensitive elements in the sentence uttered. This roughly corresponds to what Stanley and Szabo (2000, pp. 228–9) and Perry (2001, pp. 42ff.) call

<sup>16</sup> Thanks to Chris Gauker for discussion here.

the *semantic* role of context. The second way is that context plays a role in determining what is communicated by the linguistic act over and above its semantic content. This is the genuinely pragmatic role of context.

A bit later, King and Stanley distinguish between *weak pragmatic effects* and *strong pragmatic effects*, where the former involve pragmatic information that goes into the determination of the meaning and the strong pragmatic effects help me to work out what was meant based on contextual information and what was literally said (the semantic contribution). Thus strong pragmatic effects take us from our understanding of an utterance that is literally about Bob's handwriting and punctuality to our understanding that the letter writer does not think well of Bob.

I believe that a number of pragmatic effects play a role in how word meanings are modulated within a conversational context. Does this count as a strong pragmatic effect or is it a weak pragmatic effect? I'm not sure. For purposes of this monograph I am neutral on the matter.

Still, it would be nice to have neutral terminology to describe the difference between pragmatic processes that figure in the modulation of word meaning and those that figure in, for example, speech acts and conversational implicature. One thing we could do is distinguish effects that are pre-semantic from those effects that are post-semantic. Alternatively, we could take a leaf from Korta and Perry (2011) and distinguish between *near-side* pragmatics and *far-side* pragmatics. In this instance, the mechanisms by which we become entrained with each other on word meanings are near-side pragmatic mechanisms. I will remain neutral for now on whether they also count as instances of pragmatic intrusion.<sup>17</sup>

Before moving on, I should also point out the connection between the idea of microlanguages and a suggestion due to Donald Davidson (1986) that we don't learn languages writ large, but rather develop "passing theories" on the fly which we use to interpret our interlocutors. My approach is also very much in the spirit of Davidson's

<sup>17</sup> As King and Stanley (2005) persuasively show, even in familiar cases the question is very subtle.

proposal, although his proposal was a bit thin on detail. I have no objection to people thinking of this work as being one way of executing Davidson's idea in detail (I likewise have no objection to people thinking this is a way to execute some of the basic ideas of Relevance Theory).

Thus far I've been talking about the relation between semantics and pragmatics and haven't said much about the syntactic component of the grammar. The theory of the lexicon I am presenting here is largely neutral about the syntactic component, but I do want to lay my cards on the table about my independent commitment to a robust syntactic component, if only so I can say a bit about the kinds of syntactic constraints that the syntactic component of the grammar might put on the lexicon.

### *The Role of Syntax*

Much writing on language tends to treat linguistic competence as a unified phenomenon made possible by a single mechanism or module of human cognition. It seems more reasonable to suppose that the broad class of phenomena that we call "linguistic" or think of as having to do with "language" are supported by a combination of narrow mechanisms of the mind/brain. One such core mechanism would be what Hauser, Chomsky, and Fitch (2002) called the *FLN*, for "faculty of language narrowly construed."

By hypothesis, the FLN is a natural object that is part of our biological endowment. For example, Hauser, Chomsky, and Fitch speculate that this core linguistic faculty did not evolve gradually in response to selectional pressures, but was sudden in evolutionary terms and involved a biophysical wiring solution—a solution for hooking up the perceptual/articulatory system (the system that affords speech comprehension and production) with the conceptual/intentional system (the system that interprets and uses linguistic communication). The thesis is speculative, but not without supporting evidence. In the simplest form, support for the thesis involves the observation that low-level physical and mathematical principles underlie many of the recursive patterns that we see in nature—ranging from the spiral

patterns that we see in shells to the Fibonacci patterns we see in the distribution of seeds in a sunflower.<sup>18</sup>

To illustrate the recursiveness of natural language, consider the following very simple case.

- (1) This is the cat that ate the rat that ate the cheese that was made by the farmer that . . .

These sorts of patterns, in Chomsky's view, provide some of the evidence that the structure of the FLN is largely determined by basic biophysical properties.

Although it is an empirical question and certainly subject to revision, I believe that word meanings are *partly* determined by the FLN—that the FLN may contribute a thin rigid skeleton upon which word meanings are constructed. I also believe that the *bulk* of word meaning is determined by linguistic coordination mechanisms.

Why is it reasonable to think that the FLN constrains aspects of the lexicon? This certainly seems to be the conclusion one would draw from work by Mark Baker (1988), which argues that morphological and lexical properties are actually determined by the syntax (hence FLN).<sup>19</sup> Baker's thesis involves the way that complex syntactic principles become incorporated into lexical items, but we can also look to cases where more basic notions like thematic structure seem to be lexically encoded. Following Higginbotham (1989a) we can illustrate the basic idea by considering the following fragment from Lewis Carroll's poem "The Jabberwocky."

- (2) Twas bryllyg, and the slythy toves did gyre and gymblye in the wabe . . .

Just from the surrounding syntactic environment we can deduce quite a bit about the meaning of the term 'tove'. We know, for example, that toves are the sorts of things one can count (unlike substances like

<sup>18</sup> For a good general introduction to this type of phenomena, see Stewart (1998); thanks to Noam Chomsky (pers. comm.) for the pointer.

<sup>19</sup> Thanks to Paul Pietroski for discussion here.

water), that they are spatially located and can move and undergo state changes (unlike numbers), they are quite possibly capable of acting under their own volition. All of this is defeasible, but these are reasonable suppositions to deduce from the surrounding syntactic structure.

How much information is “hard coded” into the lexical entry and not subject to modulation (or not easily subject to modulation)? Here we can afford to be open-minded, noting that important work on the nature of the lexicon by, for example, Grimshaw (1990), Hale and Keyser (1987, 1993), Pustejovsky (1995), Nirenberg and Raskin, (1987), Pustejovsky and Bergler (1991), Boguraev and Briscoe (1989) among many others points to the idea that some thematic elements of the lexicon are relatively stable. Let’s consider a concrete example. A verb like ‘cut’ might contain the sort of information proposed by Higginbotham (1989b: 467):

CUT is a V that applies truly to situations *e*, involving a patient *y* and an agent *x* who, by means of some instrument *z*, effects in *e* a linear separation in the material integrity of *y*.

Such a lexical entry might provide a kind of skeleton that gets fleshed out via a process of entrainment with our discourse partners. Some elements of the lexical entry for ‘cut’ may be quite stable—the thematic roles of agent, patient, and theme, while others may be subject to modulation (for example, what counts as a linear separation).

There may well be other constraints on word meaning that are imposed by the surrounding semantic context. In a very interesting recent work Asher (2010) has described rules that govern the way in which word meanings get coerced. All of these approaches can and should be taken on board if they can either illuminate the part of word meaning that is stable or constrain the process by which meanings are modulated.

I want to conclude this section with a word about the acquisition of word meanings by children. Given the rapid acquisition of lexical items by children during their critical learning period (ages 1.5–6) and given their corresponding acquisition and grasp of these basic thematic relations (provided only impoverished data, no reinforcement,

etc.), it seems reasonable to speculate that these thematic relations are stable precisely because they are part of the FLN, as discussed earlier. But as Bloom (2000) has argued, all of this just gives children a first pass at understanding word meanings. To flesh things out children also need to incorporate a broad range of contextual information and real world knowledge. That is, children acquire the skeleton quickly, but it takes them years to put flesh on those lexical bones. Of course, Bloom is assuming that there is a target meaning to be learned. It is more accurate to say that children, like adults, must ultimately collaborate with their discourse partners to flesh out those word meanings, and ultimately learn how to modulate those word meanings on a case-by-case basis.

Just to be clear, in laying my cards on the table about the contribution of syntax (i.e. thematic structure) to word meanings I am not saying that such a contribution is critical to a theory of the dynamic lexicon. I just happen to independently believe in these facts about the lexicon. What I am committed to, however, is that various sub-theories do constrain the way in which word meanings are modulated—that is to say, the process is not magic, and we all do it quite efficiently. If there is an explanation that is nonmagical, it presumably involves the joint action of numerous modules very much like the FLN.

## 4

# Meaning Underdetermination, Logic, and Vagueness

In the previous three chapters we looked at how word meanings are underdetermined and subject to modulation. In this chapter our goal is to come to grips with some of the consequences that ensue once we recognize that meanings are underdetermined in this way.

Just to make the discussion vivid, I want to reiterate that this kind of underdetermination holds for possibly<sup>1</sup> every predicate that we use. I believe that this kind of meaning underdetermination even holds for cases of mathematical predicates like ‘straight line’. We might think that the notion of straight line is perfectly sharp, but new developments in mathematics can raise questions about just how sharp that notion is. In fact, they already have, thanks to the development of non-Euclidean geometries. For example, in a four-dimensional affine space a line can be perpendicular to itself. On learning this, you might ask, well then is it really a straight line? And of course, one could very well have said “no.” A decision had to be made about whether the “lines” in this new geometry fell under the meaning of ‘line’ in the old. Just like new scientific discoveries and technological developments can put our modulation of ‘person’ under pressure, so too can discoveries in mathematics put our modulation of ‘line’ in question.

<sup>1</sup> The hedge ‘possibly’ is necessary because if the meaning of ‘underdetermined’ is underdetermined then there must be cases of lexical items wherein we have not determined if they are underdetermined. Perhaps we will determine that some of them aren’t.

This point is basically just an extension of a point that has been made by Kitcher (1981) and others about the notion of formal rigor in mathematics generally. We may think we have sharpened up a notion like ‘infinitesimal’, but the operational definition of rigor is always “rigorous enough for current purposes.” One natural way of taking Kitcher’s point is that there is no absolute notion of rigorization.<sup>2</sup>

All of this leads to an interesting problem. Braun and Sider (2007) more or less agree that language is underdetermined in the way I have described,<sup>3</sup> and they believe that this underdetermination may extend to nearly all predicates but they believe that this means there is a problem with providing a traditional semantics for natural language. We take up their argument in the next section.

## 4.1 Meaning Underdetermination and Semantics

Consider the following sentences.

- (1) Kansas is flat
- (2) ‘Kansas is flat’ is true iff Kansas is flat

There are plenty of circumstances where we might say that (1) and (2) are true (or if we are unimpressed by the flatness of Kansas, we might say (1) is false), but Braun and Sider (2007) don’t think utterances of these sentences (or perhaps *any* sentences) can ever be strictly speaking true (or false). On their view the problem is that, while meanings are underdetermined, semantic objects are not:

We assume that the properties, relations, and propositions that are candidates for being the meanings of linguistic expressions are precise: any  $n$ -tuple of objects either definitely instantiates or definitely fails to instantiate a given  $n$ -place relation, and any proposition is either definitely true or definitely false. (Braun and Sider 2007: 134)

<sup>2</sup> A similar point is made by Arana and Mancosu (2012) in the context of the relation between plane and solid geometry.

<sup>3</sup> They use the term ‘vagueness’ but use it in a broad enough way so that it approximates what I mean by underdetermination. Their point carries over directly to underdetermination.

The idea behind their claim is simple enough. We suppose that the semantic values introduced by a semantic theory will look something like one of the following options, where the double vertical strokes indicate the meaning or semantic value of the enclosed expression.

$\| \text{'flat'} \| = \text{the property of being flat}$

$\| \text{'flat'} \| = \{x: x \text{ is flat}\}$

$\| \text{'flat'} \| = f: D \rightarrow \{0, 1\}$ , For all  $x \in D$ ,  $f(x) = 1$  iff  $x$  is flat

$\| \text{'flat'} \| = f: D' (D' \text{ a subset of } D) \rightarrow \{0, 1\}$ , For all  $x \in D'$ ,  $f(x) = 1$  iff  $x$  is flat

Notice that all of these are precise—even the one that introduces a partial function!<sup>4</sup> (To see this, note that for every object in the domain, it is precisely determined whether it is in the extension of set of ‘flat’, in the anti-extension of ‘flat’, or undetermined.)

If semantic values must look like this, then according to Braun and Sider it follows that almost all utterances of (1) and (2) are not, strictly speaking, true (or false). They fail to be true (false) because for them to be true (false) there would have to have precise semantic objects as their meanings. But there are no such precise semantic objects. Their conclusion:

[T]ruth is an impossible standard that we never achieve... [I]t would be pointlessly fussy to enforce this standard to the letter, requiring the (exact) truth... nor would it be desirable to try, for the difference between the legitimate disambiguations of our sentences are rarely significant to us. (2007: 135)

They opt for an alternative notion according to which sentences of natural language are not true, but are “approximately true.”

Speaking vaguely (as always), there is a range of *legitimate disambiguations* for a vague expression... When all the legitimate disambiguations of a sentence are true, call that sentence *approximately true*.

<sup>4</sup> Paul Pietroski has suggested to me that it might be possible to challenge this assumption. Given that the characterization of the property/extension/function in each case involves the use of the underdetermined term ‘flat’, one might wonder if these are precise objects after all. It is hard to work through this. Would the idea be that properties are underdetermined objects? There is room to maneuver with properties, but what about the set  $\{x: x \text{ is flat}\}$ . Does it make sense to talk about an underdetermined set (or function for that matter)? Note we don’t want to confuse this with the idea of a fuzzy set in the sense of Zadeh (1965) which is, in its own way, a fully determinate object.

So we rarely (maybe never) speak truly; we are rather saying something approximately true, which is fine according to Braun and Sider.<sup>5</sup> The view bears similarities to Lewis (1982) except that on Lewis's view when all of the legitimate disambiguations are true then the sentence is super-true (not merely approximately true, not merely true, but "super-true").<sup>6</sup>

There are of course a host of other views in the neighborhood of the Braun and Sider view, all of which endorse the basic idea that language is vague or ambiguous or indeterminate or underdetermined and that the objects of semantic theory are not. Some of the views are substantially more pessimistic than Braun and Sider—for example Unger (1979: 249) has notoriously held that "our existing expressions, at least by and large, fail to make any contact with whatever is there." I don't intend to sort out these different positions because I think all of them rest on a mistake—the mistake being the assumption that semantic objects are precise. Not only is it assumed that semantic objects are precise, but I think it is also assumed that to reject semantic precision would be to reject a viable working science.<sup>7</sup>

In the first place, semantics as practiced in linguistics departments is typically interested in the notion of *truth in a model*. Semanticists typically don't worry about whether we can identify the set of all things that are 'flat' because different models will assign different sets to that term. It is only when we try to press semantic theory into the service of delivering an absolute semantics (one for truth simpliciter) that problems

<sup>5</sup> It is, however, as they point out, only approximately true that what we say is approximately true. I'm unclear on whether they take the language of semantics to be approximately true.

<sup>6</sup> However Braun and Sider believe Lewis (1993: 29) is hedging in the following passage (their emphasis): "Super-truth, with respect to a language interpreted in an imperfectly decisive way, *replaces* truth simpliciter as the goal of a cooperative speaker attempting to impart information." I'm unclear on what the force of this point is supposed to be, given that super-truth can do all the work that truth was supposed to do.

<sup>7</sup> Teller (2012: 8) provides an example. "There is something I want to emphasize: I do not advocate giving up the familiar framework that includes presuppositions 1) [that semantic values, in particular propositions, are precise] and 2) [that for a statement to be true is for the statement to express a true proposition]. The mystique of Kuhnian paradigm shifts to the contrary, science does not generally discard frameworks that have proven broadly fruitful... [T]he framework of formal semantics will continue to provide the best way of understanding many important features of language."

arise. Of course some semanticists are interested in utilizing truth simpliciter in a semantic theory—typically, this was part of the project in natural language semantics envisioned by Higginbotham (1985) and Larson and Segal (1995). I happen to be involved in this project as well. And as Lepore (1983), Higginbotham (1985), and others have stressed, if you want a theory of meaning that connects language with the world, then truth in a model is not sufficient—we will need to deploy truth simpliciter. But is it really the case that the semantic objects in an absolute semantics must be (or even are) precise? I don't think so.

To explain this idea, I need to introduce a toy semantic theory within a truth-theoretic framework as envisioned by Davidson (1967) and developed by Higginbotham (1985), Larson and Ludlow (1993), and in considerable detail Larson and Segal (1995). On such semantic theories, instead of introducing model-theoretic objects as semantic values we will instead offer a truth-conditional semantics that makes no reference to such objects but still provides a recursive semantics that carries all of the information we are interested in for the conduct of semantic investigation.

In the fragment that follows I define a language *L* using some simple recursive rules and then provide a semantics for *L* that shows how we can compute the semantic value of the whole sentence from the semantic values of the parts. We specify the syntax of our toy language as follows:

#### Syntax of *L*<sup>8</sup>

- (i)  $S \rightarrow S_1 \text{ and } S_2$
- (ii)  $S \rightarrow S_1 \text{ or } S_2$
- (iii)  $S \rightarrow \text{it is not the case that } S_1$
- (iv)  $S \rightarrow \text{NP VP}$
- (v)  $\text{NP} \rightarrow \text{'Michael Jordan', 'Kansas'}$
- (vi)  $\text{VP} \rightarrow \text{'is flat', 'leaps'}$

<sup>8</sup> Clearly this fragment not only ignores plenty of structure (like the structure of NP and VP but also takes certain liberties in e.g. classifying 'is flat' as a VP when it is probably better classified as an adjective phrase (AP). This won't matter for our purposes. More robust fragments can be found in Larson and Segal (1995).

To illustrate, successive applications of the rules (iv), (v), and (vi) would yield the following bracketed structure:  $[_S[_{NP} \text{'Kansas'}]_{VP} \text{'is flat'}]]$ , where the syntactic categories of the constituents are represented by bracket subscripts.

For the semantics we introduce axioms for the terminal nodes (that is, for the lexical items) and additional axioms that tell us, for the various syntactic structures, how we can compute the semantic value of a constituent from the semantic values of the constituents it contains. In effect, we begin by using the axioms for the semantic values of the words, and then use the nonterminal axioms when we compute the semantic value of the whole sentence. To keep things clean and consistent with most of formal semantics, I will suppose that the semantic value of the sentence is a truth value. We also introduce the predicate *Val*, where  $Val(A, B)$  says that *A* is the semantic value of *B*.

With that, the axioms for our terminal nodes (words) will be as follows.

- (3)
- a.  $Val(x, \text{'Jordan'})$  iff  $x = \text{Jordan}$   
 $Val(x, \text{'Kansas'})$  iff  $x = \text{Kansas}$
  - b.  $Val(x, \text{'is flat'})$  iff  $x$  is flat  
 $Val(x, \text{'leaps'})$  iff  $x$  leaps

So axiom (3a) tells us that *x* is the semantic value of 'Kansas', just in case *x* is identical to Kansas. (3b) tells us that *x* is a semantic value of 'flat' just in case *x* is flat.

The nonterminal nodes now tell us how to compute the semantic value of higher level syntactic structure (like the verb phrase and the sentence as a whole) from the semantic values of the words. Since this is a very simple language we only need the following rules to cover all the cases.

- (4)
- a.  $Val(T, [_S NP VP])$  iff  
for some *x*,  $Val(x, NP)$  and  $Val(x, VP)$
  - b.  $Val(x, [_\alpha \beta])$  iff  $Val(x, \beta)$  (where  $\alpha$  ranges over categories, and  $\beta$  ranges over categories and lexical items)

- c.  $\text{Val}(T, [_{\mathcal{S}} S_1 \text{ 'and' } S_2])$  iff  
 $\text{Val}(T, S_1)$  and  $\text{Val}(T, S_2)$
- d.  $\text{Val}(T, [_{\mathcal{S}} S_1 \text{ 'or' } S_2])$  iff  
 either  $\text{Val}(T, S_1)$  or  $\text{Val}(T, S_2)$
- e.  $\text{Val}(T, [_{\mathcal{S}} \text{ 'it is not the case that' } S_1])$  iff  
 it is not the case that  $\text{Val}(T, S_1)$

Finally, we need some derivation rules that tell us how to mechanically apply the axioms. In this case we only need two rules. The first one tells us that if we derive something of the form  $\alpha$  iff  $\beta$ , we are entitled to swap in  $\beta$  for  $\alpha$  at any point in the derivation. The second rule does a simple cleanup to get rid of the variables we have introduced.

Derivation Rule (SoE)

$$\begin{array}{c} \dots \alpha \dots \\ \underline{\alpha \text{ iff } \beta} \\ \text{therefore} \dots \beta \dots \end{array}$$

Derivation Rule (SoI)

$$\begin{array}{c} \underline{\phi \text{ iff for some } x, x = \alpha \text{ and } \dots x \dots} \\ \text{therefore } \phi \text{ iff } \dots \alpha \dots \end{array}$$

Using these axioms and the two derivation rules we can, in a straightforward way, derive theorems like (2). Let's walk through the derivation for 'Kansas is flat'.

- (i)  $\text{Val}(T, [_{\mathcal{S}_2} [_{\text{NP}} \text{ 'Kansas' } ] [_{\text{VP}} \text{ 'is flat' } ]])$  iff  
 for some  $x$ ,  $\text{Val}(x, [_{\text{NP}} \text{ 'Kansas' }])$  and  $\text{Val}(x, [_{\text{VP}} \text{ 'is flat' }])$   
 [instance of axiom 4a]
- (ii-a)  $\text{Val}(x, [_{\text{NP}} \text{ 'Kansas' }])$  iff  $\text{Val}(x, \text{ 'Kansas' })$
- (ii-b)  $\text{Val}(x, [_{\text{VP}} \text{ 'is flat' }])$  iff  $\text{Val}(x, \text{ 'is flat' })$   
 [both instances of (4b)]
- (iii)  $\text{Val}(x, \text{ 'Kansas' })$  and  $\text{Val}(x, \text{ 'is flat' })$   
 [from steps (i) and (ii-a), (ii-b) by applications derivation rule SoE]

## 108 UNDERDETERMINATION, LOGIC, AND VAGUENESS

- (iv) for some  $x$ ,  $x = \text{Kansas}$  and  $x$  is flat  
[from step (iii) and axioms (3a) and (3b) by applications of SoE]
- (v) Kansas is flat  
[from (iv) by application of derivation rule SoI]

So we have derived the following: [<sub>S</sub> [<sub>NP</sub> ‘Kansas’] [<sub>VP</sub> ‘is flat’]] is true iff Kansas is flat.

Our fragment is quite simple but, as demonstrated in Larson and Segal (1995), it will scale up nicely with the addition of more axioms. Indeed, these types of fragments can be extended to cover the same data covered by a traditional model-theoretic semantics. That is to say, if you chose to, you could do semantics in this way without losing anything, modulo some story about representing entailment relations.<sup>9</sup>

But now notice that we did all this without introducing the usual machinery (utilizing properties, sets, functions) that is supposed to be precise (and determinate). What is interesting about our axiom for ‘flat’ is that it takes advantage of the underdetermined meaning of ‘flat’ in the metalanguage. Notice that there is no barrier to using axioms like this in a meaning theory in a way that will deliver theorems like (2), repeated here.

- (2) ‘Kansas is flat’ is true iff Kansas is flat

In fact we just proved that we can derive (2) without the technology of sets, properties, relations, etc.

Let me make it clear that I am not making a claim against contemporary formal semantics. There is no question but that formal semantics has been extremely fruitful and that it has provided many profound discoveries and insights into the nature of language. In the past forty years we have seen important discoveries in generalized quantifier theory, in the way that adjectival modification works, in

<sup>9</sup> For this you might retain a traditional model-theoretic semantics, or you could run your theory of entailment off of the syntax as e.g. in Ludlow (1996, 2002). Gauker (2013) suggests relying on a notion of truth in a context, which we could redraw to something like truth in a microlanguage. As we will see, there are some complications to deal with.

the way that pronominal anaphora works. We have seen advances in the theories of tense, aspect, modality, and the list could go on and on. None of these discoveries, however, rests upon the assumption of a precise or determinate metalanguage, although, as we will see, perhaps some intractable semantic puzzles (like vagueness) do rest on that assumption.

The point is that for the most part all the actual results of semantics are preserved if we reject the idea of precision in the metalanguage. So why do people suppose the metalanguage must be precise?

Concern about vague and underdetermined objects isn't the only concern driving the belief that we need to introduce precise or well-defined objects into our semantic theory. Here we are getting into the area of the sociology of philosophy, but it might be useful to talk about it since I believe that most of the pressure for semantic precision is born of old habits and unchallenged assumptions. In particular I think that the assumption is the residue of two hangovers—the Quinean hangover and the Platonic hangover.

Quine's project, of course, was to construct a language, or more accurately, to *regiment* natural language into a form that was optimal for the conduct of scientific inquiry. For Quine, this meant that the language needed to be extensional, first order, etc. If you extend this idea to semantics then you are effectively saying that if you want to carry out a project of scientific semantics then you have to do so in an extensional metalanguage.

Davidson, of course, was explicit about the extensional constraint for semantics—he was in effect making the case for a semantics that was light on ideology.<sup>10</sup> But there really is no reason to think semantics must take place in an extensional metalanguage. That is to say, we can give a semantics for natural language in a metalanguage that contains the predicates 'necessarily' and 'possibly'.

<sup>10</sup> Although, I should note that Davidson's position on underdetermination is equivocal. Ernie Lepore (pers. comm.) has a reading of Davidson (1967) in which Davidson was proposing that we could lift not only ambiguity but meaning underdetermination (in my sense) into the metalanguage. I'm not able to form a strong view about Davidson's position here.

Likewise, many semanticists have held that we need to give our semantics for tense in a detensed metalanguage. It is fine to have primitives like earlier-than and later-than, but it would be error to think that we can give a semantics for tense in a tensed metalanguage. But why? We give our semantics for quantifiers in a language with quantification, and we give our semantics for negation in a language with negation. Why should tense be special? Why must tense be given in a metalanguage without tense?

As it turns out, one *can* give a reasonable answer to the question about tense; you can argue that a tensed metalanguage will yield McTaggart-type paradoxes. I have my doubts about whether it does lead to this, but at least it is an answer. It is harder to see why a metalanguage could not have underdetermined meanings.

In point of fact we don't conduct science in a special language, and whatever precision we need in science is precision of method or instrumentation that we introduce explicitly as theoretical machinery. We don't rely on our language of inquiry to provide this precision.

No one in physics or chemistry, so far as I know, thinks that they need a special metalanguage in which to conduct scientific inquiry,<sup>11</sup> so why should semanticists be expected to do so? This is the Quinean hangover. We *think* that semantics requires a regimented metalanguage, but we don't really reflect on *why* we suppose this. It really appears to just be an ingrained habit that comes down from Quinean admonitions about the proper way to conduct scientific inquiry.

Now obviously, the Quinean hangover is really the residue of some other hangovers, so maybe we can push the problem back to Carnap or further. My point here is not to assess blame, but just to observe that the need for absolute precision is an unargued assumption about the nature of semantics.

If we really push back on the blame game, however, we get to what I call the Platonic hangover. Here the idea is that precision involves approximation to some target. So, for example, take a concept like justice. As we move through Plato's dialogue *The Republic* we get a sharper

<sup>11</sup> Woodger's (1937, 1939) project for biology would be a notable rare exception.

and sharper picture of what justice must be. But Plato seems to think that we are getting a sharper and sharper picture by virtue of our getting closer and closer to the *form* of justice (strictly speaking we are starting to recall the form with more clarity). We have long since abandoned Platonic forms, but there is a residue that we hang onto—we cling to the idea that there is a target—an exact precise notion of justice (the *meaning* of the term ‘justice’).

I agree that Plato is elucidating a number of important concepts and they are getting more and more precise, but I don’t agree that this is because we are getting closer to the concepts themselves as they rest in Plato’s heaven. I would argue that we are merely coming up with better and better modulations—or if you prefer, we are constructing better and better concepts. What makes them better is not that they are closer to some perfect target, but rather that (as we saw in Chapter 3) we are coming up with progressively more serviceable modulations via a normatively constrained process of argumentation.

Notice that this same Platonic hangover seems to be at play in our discussion of semantics when we assume that there must be some privileged meaning to a term like ‘flat’ or ‘knows’, and in particular when we suppose that this privileged meaning must have a kind of perfection built into it—flat things have to be utterly free of surface irregularities. Knowledge has to be absolutely immune to error. Those may be possible modulations of the word meanings, but there is no reason to privilege them.<sup>12</sup> We should instead embrace the doctrine that I earlier called Meaning Egalitarianism.

Are we home free then? Not yet, because Braun (pers. comm.) has a rejoinder. We can set it up with our examples (1) and (2), repeated here.

- (1) Kansas is flat
- (2) ‘Kansas is flat’ is true iff Kansas is flat

Now let’s suppose that I’m conflicted about whether Kansas is flat (I’ve been there and it is definitely hillier than Southwestern Minnesota). If

<sup>12</sup> Here is the point where my view of modulation breaks from Recanati’s, for he does seem to take the absolutely sharp meaning as being privileged, and we modulate from that.

I say that (1) is neither true nor false, then I have to say that (2) is neither true nor false. For if (1) is neither true nor false, then the left-hand side of (2) is false, while its right-hand side is neither true nor false. In fact, I will have to reject all instances of the T-schema (T) in which S is neither true nor false.

(T) ‘S’ is true iff S

Braun is raising a good point, because it raises the question of whether I intend to give up bivalence, and if so, what the consequences are for instances of the T-schema. I think it would be interesting to pursue that option (giving up bivalence), but I have a different story to tell.

To set up my story I first want to make it clear that I take the semantic theory to be a theory which computes the semantic values of utterances (or, if you prefer, tokens)—not sentences construed as abstract objects (this is a distinctively anti-Kaplanian assumption which I won’t defend here).

On my view, in any microlanguage, admissible *utterances* having the form of (1) or (2) must be either true or false. How is this possible, given that meanings are underdetermined? Let’s return to my claim that when we engage in conversation we build microlanguages on the fly. Not only are word meanings dynamically narrowed or broadened in these cases, but there is also a question of which words (and hence sentences) make it into the microlanguage. My view is that no sentence is admissible until it is sharp enough for us to assert a claim that is clearly either true or false. Let’s state this as a principle.

*Microlanguage S-admissibility:* No utterance *u* of a sentence *S* is admissible in a microlanguage *L*, unless discourse participants (tacitly) agree that the terms of *S* are modulated so that an utterance of the sentence will be determinably either true or false.

Now that may seem like a bizarre claim, since, after all, we routinely do say things that are supposed to have indeterminate truth values. We give examples in class like ‘Ludlow is bald’, after all.

But there is a difference between talking about sentences with indeterminate truth values and introducing them into a conversation as

vehicles for assertions. I would say that, when we talk about sentences with indeterminate truth values, we are talking about whether utterances of them should be introduced into the microlanguage. When we are saying the truth value is indeterminate, we are saying that it is not admissible—it cannot be deployed to make an assertion.

Of course, sometimes people say things like ‘Ludlow is bald’ and the listeners may nod even though they may have previously thought it is indeterminate whether I am bald. To return to an example I gave in Chapter 1, Chris Barker (2002) has made the observation that this often is not merely making a claim about my baldness, but it is actually an instruction about where to push the meaning of ‘bald’. It is saying that (or offering that) ‘Ludlow is bald’ is a safe application of ‘bald’—hereafter, everyone with less hair than Ludlow (modulo head size) is safely bald.<sup>13</sup>

All of this raises interesting questions about how we are to think of Semantic Theory (in capital letters to indicate that I am talking about the scientific enterprise of semantics). We normally think of a semantics like the toy version I gave above as something that works for a stable language. Maybe the theory is something that a third party deploys for someone else’s language, or maybe it is used in the interpretation of our own language, but the question is, what is the nature of a toy semantics like the one I gave above, given our understanding of microlanguages and the microlanguage admissibility constraint?<sup>14</sup>

The answer is that the toy semantics is a one-off “passing theory” for computing the semantic values of expressions in a particular microlanguage L (I believe, based on personal communication, that this was Davidson’s view as well—one-off T-theories are used to interpret the

<sup>13</sup> Parenthetically, I think that Barker missteps when he says that “it is the very act of using a vague expression that makes it more precise.” Saying “Ludlow is bald” doesn’t make ‘bald’ more precise because it doesn’t speak at all to the edge cases. It is saying “Ludlow is safely bald for our purposes.” I would prefer to say that utterances like this broaden the meaning of a term because they introduce more candidates as semantic values of a term like ‘bald’. Paradoxically, while nailing down one specific case of baldness, it is not necessarily making the term itself sharper. The only way to do that would be to tell us about the edge cases, but once accepted, ‘Ludlow is bald’ is no longer an edge case. It is a safe case for purposes of our conversation.

<sup>14</sup> I am indebted to John MacFarlane for conversations here.

passing theories like those introduced in his “Derangement” paper). On the assumption that Semantic Theory is interested in the psychological mechanisms (clearly no longer a Davidsonian assumption) by which we assign meanings (semantic values) to linguistic forms, then Semantic Theory is interested in understanding the underlying system that allows us to build such semantic theories on the fly. Put another way, we can think of ourselves as building little toy grammars for fragments of “languages” all the time, complete with semantic theories for interpreting them. Semantic Theory is the study of how we do this.

Presumably, there must be stable elements to the construction of a passing theory; otherwise it is difficult to see how it could be possible. It is a safe bet that most of the nonterminal semantic rules are stable across these shifts (they may well be stable across all human languages). Thus the real dynamic portion would be in constructing the terminal (lexical) rules on the fly.

## 4.2 Word Meaning Modulation and Logic

As it turns out, the dynamic lexicon not only shakes up some of our assumptions about semantics, but it also shakes up our understanding of what logic looks like. On the face of it, if *every* term is underdetermined and dynamic, then in any argument that is presented in time (as natural language arguments are) we cannot guarantee that equivocation is not taking place. To see the problem consider the most trivial possible logical argument:

F(a)

Therefore F(a)

If the meaning of F shifts, the argument may not be sound even if Fa is true—a kind of equivocation might have taken place. Does this mean that logic goes out the window? Not at all. For expository purposes let’s sequentially number occurrences of terms in an argument, so that, for example, in the argument we just gave the form is the following.

F<sub>1</sub>(a)

Therefore F<sub>2</sub>(a)

Again, we are saying that the term is  $F$ , and that  $F_1$  and  $F_2$  are *occurrences* of the term  $F$  within the argument. To keep our understanding of validity stable, let's say that soundness calls for a third constraint in addition to the validity of the argument and the truth of the premises.

It appears that the argument is sound if the meaning is stable between  $F_1$  and  $F_2$  but also if  $F_2$  is a broadening of  $F_1$  (a narrowing or a lateral shift in meaning will not preserve truth). Let's take a concrete example.

Jones is an athlete<sub>1</sub>

Therefore Jones is an athlete<sub>2</sub>

If a shift has taken place between premise and conclusion (between the meaning of 'athlete<sub>1</sub>' and 'athlete<sub>2</sub>') it cannot be a shift that rules out individuals that were recognized semantic values of 'athlete<sub>1</sub>'. If 'athlete<sub>1</sub>' admits racecar drivers and 'athlete<sub>2</sub>' does not, then the argument is not valid. If the second occurrence broadens the meaning of the term 'athlete', the argument *is* sound.

Broadening meaning doesn't always ensure validity. Let's introduce a negation to the argument form we just gave.

Jones is not an athlete<sub>1</sub>

Therefore Jones is not an athlete<sub>2</sub>

This time matters are reversed. Assuming the premise is true, the argument is valid just in case either 'athlete<sub>2</sub>' preserves the meaning of 'athlete<sub>1</sub>' or it is a *narrowing* from 'athlete<sub>1</sub>'. Negation isn't the only environment that dislikes the broadening of meanings. Consider the following.

If Jones is an athlete<sub>1</sub> then Jones is healthy<sub>1</sub>

Jones is an athlete<sub>2</sub>

Therefore Jones is healthy<sub>2</sub>

For the argument to be valid 'athlete<sub>1</sub>' can be broader than 'athlete<sub>2</sub>', but it cannot be narrower. 'healthy<sub>2</sub>' can be broader than 'healthy<sub>1</sub>' but not narrower. Notice that this seems to hold if we reverse the order of the premises as well.

Jones is an athlete<sub>1</sub>  
If Jones is an athlete<sub>2</sub> then Jones is healthy<sub>1</sub>  
 Therefore Jones is healthy<sub>2</sub>

Again, ‘athlete<sub>1</sub>’ can be narrower than ‘athlete<sub>2</sub>’ but it cannot be broader.

What is going on here? Is there a way to make these substitution rules systematic within natural language? I believe that there is, because I believe that they track what linguists call upward and downward entailing environments. To a first approximation, an upward entailing environment is one where a predicate with a broader range can be swapped in for a predicate with a narrow range and truth is preserved; a downward entailing environment is one where any predicate with a narrower range can be swapped in for a predicate with a broader range and truth is preserved. Elsewhere (Ludlow 1995, 2002) I have argued that these environments can be syntactically identified.

Let’s call an occurrence of a term in an upward entailing environment a *positive occurrence* of the term, and let’s call an occurrence of a term in a downward entailing environment a *negative occurrence* of the term. Assuming that we can identify these environments, we can state a constraint on valid arguments as follows:

*Dynamic Lexicon Constraint on Validity (DLCV)*: if *t* is a term with multiple occurrences in an argument and it plays a direct role in the derivation of the conclusion,<sup>15</sup> then those occurrences must either

<sup>15</sup> The proviso that it “plays a direct role in the derivation of the conclusion” is designed to allow us to ignore terms that have multiple occurrences but that are inert in the structure of the argument. (I am indebted to an anonymous reviewer for pointing out the need for this proviso.) So e.g. consider the following argument form:

$$\begin{array}{l} (P \vee Q) \rightarrow R \\ Q \rightarrow S \\ \frac{P}{R} \end{array}$$

The terms in *Q* have multiple occurrences, but they do not play a role in the derivation of *R*. Matters are a bit more complicated than this example lets on, however, because sometimes there are multiple ways to derive the conclusion in a sound argument. E.g.:

$$\begin{array}{l} (P \vee Q) \rightarrow R \\ Q \\ \frac{P}{R} \end{array}$$

have the same meanings, or be broadenings/narrowings of each other as follows:

- (i) If a term  $t$  has an occurrence  $t_1$  in the premises and  $t_2$  in the conclusion, then
  - if  $t_2$  has a positive occurrence it must be broader than  $t_1$
  - if  $t_2$  has a negative occurrence, it must be narrower than  $t_1$
- (ii) If a term  $t$  has two occurrences in the premises of the argument (i.e. in a two step chain of argumentation), then the positive occurrence must be narrower than the negative occurrence.

This constraint needs to be generalized and a proof is called for, but we can see that it works for familiar cases like arguments involving *modus ponens* (as above) and the Aristotelian syllogism. Consider the Barbara schema for example.

All dogs are things that bark  
All collies are dogs  
 All collies are things that bark

Let's annotate the terms with polarity markers + and - to indicate positive and negative occurrences as they are traditionally assigned.

All dogs<sub>1</sub>[-] are things that bark<sub>1</sub>[+]  
All collies<sub>1</sub>[-] are dogs<sub>2</sub>[+]  
 All collies<sub>2</sub>[-] are things that bark<sub>2</sub>[+]

Our Dynamic Lexical Constraint on Validity tells us that the argument is only valid if 'collies' is stable or narrows, 'barks' is stable or broadens (by i), and 'dogs' is stable or narrows (by ii). This is clearly correct. I leave it as an exercise for the reader to examine the other forms of the syllogism.

Is this a hack? To the contrary it dovetails very nicely with some deep insights that have been made into the nature of logical

We can either define soundness relative to a derivation (so that derivations, rather than arguments, are considered sound), or we can say that the form of the argument is sound if there is at least one derivation path to the conclusion such that the argument consisting of just that derivation path respects the DLCV. Other solutions are, most likely, available. See van Deemter and Peters (1997) for a survey of articles on the general problem developing logics that tolerate ambiguity, indeterminacy, and by extension underdetermination.

inference, particularly as it relates to the role of upward and downward entailing environments. Because it is important to me that this constraint be natural, I pause for an interlude on this topic. To make the point vivid (and somewhat more accessible) I'll initially make the point with Aristotelian logic and a smattering of propositional logic.

Recall that the Aristotelian syllogism stipulates seventeen valid forms. Examples, of course, include the forms in (5) and (6), both of which exemplify the form known as Barbara.

(5)

All men are humans

All humans are mortals

All men are mortals

(6)

All As are Bs

All Bs are Cs

All As are Cs

Although it's widely supposed the medieval logicians made few contributions to logic (Kant allegedly advanced this view) the medieval logicians did realize that Aristotelian logic was both too constrained and too *ad hoc*, and they sought to rectify the problem.

Aristotelian logic was *ad hoc*, because the seventeen valid forms were simply stipulated without much motivation, except that they tracked our judgments of validity. At the same time it was too constrained because it was limited to sentences that were categorical propositions; namely the following:

- All A is B
- No A is B
- Some A is B
- Some A is not B

As a result there were a number of intuitively valid inferences that fell outside of traditional Aristotelian logic, examples including (7) and (8).

(7)  
 Every man is mortal  
 Socrates is a man  
 Socrates is mortal

(8)  
 No man flies  
 Some bird flies  
 Some bird is not a man

It is of course a standard exercise of undergraduate logic texts to force natural language sentences into categorical propositions, so many of us have made or at least run across examples like the following.

(9) Socrates is a man = Every Socrates is a man  
 (10) No man flies = No man is a thing that flies

Even logical connectives present difficulties (see Sommers's (1970) treatment of sentential connectives in Aristotelian logic and discussion in Horn (1989)). Does a good Aristotelian really want to make substitutions like the following?

(11) The Yankees will win or the Red Sox will win  
 (11') All [non-(the Yankees will win) isn't [non-(the Red Sox will win)]]  
 (12) The Yankees will win and the Mets will win  
 (12') Some [The Yankees will win] is [the Mets will win]

Worries about this problem persisted at least until De Morgan (1847), who introduced the following famous case:

(13)  
Every horse is an animal \_\_\_\_\_  
 Every head of a horse is a head of an animal

The goal for many of the medievals was to expand the range of logic to cover these cases, but also to make it less *ad hoc*—that is, to avoid having to state seventeen (or more) distinct valid inference

patterns. The deep insight of the medievals was that you could perhaps reduce all of logic down to two basic rules—the *dictum de omni* and the *dictum de nullo*—with each rule to be used in a specified syntactic environment. The rules can be summed up in the following way: In the *de omni* environment (to a first approximation a positive polarity environment) you can substitute something broader in meaning (for example, moving from species to genus) and in a *de nullo* environment (to a first approximation a negative polarity environment) you could substitute something narrower in meaning (for example, from genus to species).

We can describe the general idea like this. ‘A < B’ is a general way of indicating that all As are Bs, or that every A is (a) B, or that all A is B (in effect moving from a narrower to a broader range). I’ll call the rule *Holy Grail*, because working out the details of this was in some ways the Holy Grail for medieval logic.

(*Holy Grail*)

An environment  $\alpha$  in a sentence  $\phi$  is a *dictum de omni* environment iff,

$[\phi \dots [\alpha \dots A \dots] \dots]$  entails  $[\phi \dots [\alpha \dots B \dots] \dots]$  if  $A < B$

An environment  $\alpha$  in a sentence  $\phi$  is a *dictum de nullo* environment iff,

$[\phi \dots [\alpha \dots B \dots] \dots]$  entails  $[\phi \dots [\alpha \dots A \dots] \dots]$  if  $A < B$

To see how this works, let’s return to our problematic examples from earlier. If we assume that the second (B) position in ‘Every A is a B’ is a *dictum de omni* environment, then we can simply swap ‘mortal’ for ‘animal’ in the argument below, since, following the second premise, animal < mortal. (I’ve placed the term to be substituted in bold.)

(14)

Every man is an **animal**

Every animal is mortal (animal < mortal)

Every man is mortal

But if we add a negation, as in the following argument, then we have to use the *de nullo* rule. Notice that here we can substitute ‘animal’ for

‘mortal’ (following the *de nullo* paradigm where we go from broader to narrower.)

(15)

Some man is not **mortal**

Every animal is mortal (animal < mortal)

Some man is not (an) animal

These two paradigms cover the Aristotelian syllogisms, but they cover a number of other apparently valid inferences that fall outside of the Aristotelian paradigm as well. It works just as well for the following arguments.

(16)

Every man is mortal (man < mortal)

Socrates is a man

Socrates is mortal

(17)

No man **flies**

Every bird flies (bird < flies)

No man (is a) bird

Here, since ‘flies’ is in a *de nullo* environment, we can substitute ‘bird’ for ‘flies’.

If we extend ‘A < B’ to include ‘If A then B’, then *modus ponens* is simply a specific instance of the *dictum de omni* rule (one where  $\alpha$  contains only A):

(18)

If Smith is tall then Jones is short

Smith is tall

Jones is short

And *modus tollens* is simply an instance of the *dictum de nullo* rule (again, the instance where  $\alpha$  contains only A).

(19)

If Smith is tall then Jones is short

Jones is not short

Smith is not tall

The solution to De Morgan’s example would be as in (20):

- (20)
- Every man is an animal (man < animal)
- Every head of a man is a head of a man
- Every head of a man is a head of an animal

By the 13th century Peter of Spain (1972), William of Sherwood, and Lambert of Auxere (1971) were on board with this project. In the 14th century Ockham (1951: 362) lent his support, stating that the “dicta directly or indirectly govern all the valid syllogisms.” According to Ashworth (1974: 232) by 16th century the “*Dici* (or *dictum*) *de omni* and *dici de nullo* were the two regulative principles to which every author appealed in his account of the syllogism.” (See also Sánchez Valencia (1994) for a survey of literature on this topic.)

As Zivanovic (2002) has shown, the general idea even extends to logics with the expressive power of hereditarily finite set theory, and it has been extended to certain classes of infinitary languages. So at a minimum it is a very broad-based phenomenon in logic, spanning centuries of research, and finding uptake in contemporary philosophical logic. And the Dynamic Lexical Constraint on Validity may well turn out to be a special case of the general rule of inference (Holy Grail) that I stated for the *dici de omni et nullo* paradigm.

My point here is simply that is the DLCV is not a hack but deeply motivated. The constraints on logic required by the dynamic lexicon are of a piece with the basic guiding principle of the Holy Grail of natural logic.

### 4.3 Vagueness<sup>16</sup>

As noted throughout this book, I consider vagueness to be a special case of meaning underdetermination. Not all cases of meaning

<sup>16</sup> Brian McLaughlin recently brought Shapiro (2006) to my attention, and that work is at least similar in spirit to the proposal in this section. Insofar as I can tell, the idea of meanings being open textured is similar to my notion of meaning underdetermination. There are also similarities in the idea that this basic idea can be applied to the problem of vagueness. There are, of course, significant differences in the execution of this basic idea. I’ve decided not to engage Shapiro’s proposal here because this book was basically complete when Shapiro’s work came to my attention and because my goal is principally

underdetermination count as vagueness (consider the question of whether Secretariat is an athlete, as discussed earlier) but all cases of vagueness are cases of meaning underdetermination. More precisely, cases of vagueness are those cases of meaning underdetermination that rest on (at least one) scalar dimension. Vagueness isn't a threat to bivalence because we typically don't admit excessively underdetermined expressions into our microlanguages, although they may play a role in our metalinguistic discussions about microlanguage admissibility.

What then are we to say of the Sorites argument? Recall the structure of that argument.

Having 0 hairs is bald

If 0 having hairs is bald then having 1 hair is bald

If 1 having hair is bald then having 2 hairs is bald

If 2 having hairs is bald then having 3 hairs is bald

...

If having 999,999 hairs is bald then having 1,000,000 hairs is bald

-----

Having 1,000,000 hairs is bald

The argument appears to be valid on the usual understanding of validity, and it looks as though all the premises are true (although this is disputed by many parties to the discussion about vagueness). That leaves the question of whether it respects the Dynamic Lexical Constraint on Soundness, and here we see the source of the problem.

On the dynamic conception of the lexicon, the meaning of 'bald' is shifting throughout the Sorites argument (in this respect it is similar to the "shifting sands" accounts of vagueness due to Fara (2000), Soames (1999), and Raffman (1996)). Indeed, as we noted earlier (following Barker 2002) utterances of 'x is bald' can modulate our understanding of what counts as being in the range of 'bald'. In the case of the Sorites argument here such utterances broaden our understanding of 'bald' as we proceed through the steps of the argument. It doesn't fix or

to illustrate the fruitfulness of the idea of meaning underdetermination. I would be quite pleased if there are multiple ways (or better ways) to accomplish this.

sharpen the edges—it says nothing about the edges—but it does introduce more and more elements into the range of ‘bald’.<sup>17</sup>

If this is what is going on then there is a sense in which there is no immediate puzzle here. Obviously ‘bald’ has been broadened to a ridiculous level, but there is nothing wrong with that. Let’s now add a premise to the argument to the effect that having 1,000,000 hairs is not bald. In this case we are alleged to derive a contradiction, as follows.

Having 0 hairs is bald

Having 1,000,000 hairs isn’t bald

If having 0 hairs is bald then having 1 hair is bald

If having 1 hair is bald then having 2 hairs is bald

If having 2 hairs is bald then having 3 hairs is bald

...

If having 999,999 hairs is bald then having 1,000,000 hairs is bald

-----  
Having 1,000,000 hairs is bald and having 1,000,000 hairs isn’t bald

But the Dynamic Lexical Constraints on Validity (introduced in the previous section) places demands on our understanding of the two occurrences of ‘bald’ within the conclusion. The first occurrence of ‘bald’ in the conclusion is positive, so it must be the broadest modulation of ‘bald’ in the argument. The second occurrence is negative, so it must be the narrowest modulation (possibly the one we had when starting in the first premise). So the conclusion says that having 1,000,000 hairs is bald on the broadest understanding of ‘bald’ and having 1,000,000 hairs isn’t bald on the narrowest understanding of ‘bald’—the one we had when we said that 0 hairs is bald. If this is right then the appearance of a contradictory conclusion is illusory. The argument is valid.

<sup>17</sup> If you wish, you can (following recent work in dynamic logic) treat each of the conditional premises as instructions for updating the common ground. In this case the argument wouldn’t really be an argument as above, but rather a series of instructions on modulating the meaning of ‘bald’ in the common ground. I remain officially neutral on this way of executing the idea.

One might object that there is no reason that we *must* modulate in an argument—that the argument is likewise valid if the meanings of the terms stay fixed at each step in the Sorites, but one needs to step carefully here. To see this, first suppose we take the conditionals to be material conditionals. Suppose further that the conditionals are asserted in a particular microlanguage. But by the Microlanguage Admissibility Constraint, the terms deployed must have been modulated so that what is said is clearly either true or false, so both the antecedent and the consequent in each premise must have truth values. But if they have truth values then it is safely determinate whether or not  $x$  is in the range of ‘heap’ at each step in the Sorites. If the base premises are both true and the conditional premises are all sharpened enough to be either true or false, then the modulation must be sharpened to the point that having  $n$  hairs is bald and having  $n+1$  hairs is not bald, so one of the conditional premises must be false.

Could someone dig in and say the following?: “No look, there is a fixed modulation on which all the premises are true!” They could say that, but then we can rightly ask for them to specify the modulation on which all the premises can be true. Clearly there is no *fixed* meaning on which this is possible.

If we consider strict conditionals then the same thing holds; in each world where we evaluate the material conditional, if the modulation is held fixed then one of the premises must be false in that world.

There are many other accounts of vagueness that claim these premises are false; what we have now is an explanation of why those premises sounded good. Statements of the form ‘if  $x$  is  $P$  then  $y$  is  $P$ ’ are typically metalinguistic devices for encouraging us to modulate (or not modulate) a term in a particular way (e.g. if Pluto is in the range of ‘planet’ then so must be many other Kuiper Belt objects). We don’t balk at these statements when we hear them because they are perfectly reasonable tools for persuading us to modulate word meanings. We should not let the reasonableness of these statements induce our acceptance of stipulated non-metalinguistic versions, however.

By the way, the theory just developed has some advantages over the traditional shifting sands accounts of modality. Consider Delia Fara’s

work on vagueness, beginning with her (2000) “Shifting Sands” paper and her subsequent development of that theory. On this account, the way we avoid the Sorites paradox is that the meaning of a term like ‘heap’ or ‘bald’ can shift as we march through the Sorites. Actually, the way I put the point is not strictly speaking correct for Fara—she actually holds the core meaning of these predicates constant but has them relativized for interests. The idea being that as we move through the Sorites our interests shift. Since ‘heap’ is relativized to the speaker’s interests, its extension changes as the speaker’s interests change while moving through the Sorites.

The intended advantage of Fara’s view is that it circumvents problems that infect other shifting sands versions. So, for example, consider a proposal in which you shift the meaning of the word or introduce a parameter for the standard of baldness or heapness. Now suppose you were forced to move through the Sorites in the following fashion, “that is a heap, and that is too, and that is too . . .”. On meaning shift accounts and implicit parameter accounts, the parameter for interests should be fixed in the very first iteration of ‘heap’. Thus the sands should not be able to shift as we move through the series of conjunctions (expressed “and that is too”). What is supposed to make Fara’s proposal superior to the others is that the ellipsed VP, when reconstructed, yields another interpretation of the same interest-relative predicate (because the shifting interests shift the interpretation). Actually, however, this is overly optimistic, as the relativization to speakers’ interests should be fixed in the first utterance, and remain fixed through each subsequent utterance in which the VP is ellipsed (even though the speaker’s interests may be shifting in the process).

There are other worries about Fara’s solution. Stanley (2003) objects that the modal profile of sentences like ‘that is a heap’ go badly on the Fara analysis. In particular he asks that we consider worlds in which there are no people and hence no interests. How do we evaluate an index for personal interests in those other worlds? Fara (2008) responds that she can rigidify the notion of interest to the actual world, hence, when you evaluate ‘that is a heap’ in some other possible world you bring along your interests from this world. That’s a reasonable

move,<sup>18</sup> but there is another, somewhat simpler response if we opt for an account in which the conditionals in the Sorites are simply conditionals of metalinguistic persuasion.

We don't need to index a property to interests; we can easily keep introducing the same predicate, because the predicate is simply 'x is in the range of 'heap' '. The ellipsed VP in each instance of 'so that is too' would be 'is in the range of 'heap' too'. As we march through the Sorites, when we unpack the ellipsed VP the resulting conditional of metalinguistic persuasion keeps broadening the meaning of 'heap'.<sup>19</sup>

One might think that the dynamic lexicon account comes in for its own modal profile objection, however. I've said that word meanings are underdetermined, but suppose someone really pressed on this and said that, no, actually even the dynamic theory is forced into accepting a definitive extension for each use of the term 'heap'. Why? Because even though no one in fact reasons analogically from canonical cases of something in the range of a predicate, for example 'planet' or 'heap', to *everything* that would end up in the range, we can ask what would be in the range if they *had* done this. That is, if we want to know the actual extension of a term that I use, we simply consider the worlds in which I did determine whether something was in the range of the predicate. We take the answers from all the closest worlds where I make the determination, and what we get is a fixed extension. So meaning isn't underdetermined after all. Perhaps it is just underspecified. Or so goes the argument.

But again, we need to move carefully here. In the case of a predicate like 'heap', we are in effect being asked to consider, for every object x, the following.

<sup>18</sup> But see Armstrong and Stanley (2011) for criticism.

<sup>19</sup> You might argue that the predicate 'in the range of 'heap' ' should be fixed in the very first step of the series of ellipsed Sorites statements. Let's suppose that is right. There is still a bit of normative content to these conditionals that we can hear as 'should be in the range of 'heap' '. If that's right then the full predicate can remain constant through each iteration of the ellipsed Sorites statement. There need be no shift in meaning. At the end of the series the listener can either take on the proposed modulation or reject it or accept some unspecified portion of it—allowing that only some cases should clearly be 'in the range of 'heap' '.

## 128 UNDERDETERMINATION, LOGIC, AND VAGUENESS

- (21) If I had applied sufficient analogical reasoning to *x*, I would have determined if *x* was in the range of ‘heap’, or not, or if I was undecided.

The theory is that if we do this, the result is sharp boundaries—we end up fixing the extension, anti-extension, and the undefined instances of ‘heap’.

The problem is that this doesn’t work, and it doesn’t work because the expression ‘would have’ in our counterfactual is underdetermined. As long as it remains underdetermined we can’t use this device to fix an extension.

But can’t we just skip the English version of the counterfactual and directly consider the possible worlds in which I apply analogical reasoning? That doesn’t work either, because the conditional is prior to the possible world analysis of conditionals. In fact, the introduction of possible worlds to account for the conditional fails for precisely the same reason that the introduction of sets and properties failed in more basic constructions. The introduction of precise semantic objects yields models that misfire on precisely the features of language that are important to us here—the fact that, for example, meaning is underdetermined. If you believe in meaning underdetermination then you can and should be cautious about the move to precise mathematical machinery in the semantics because, if the artifacts (like precise extensions or functions from possible worlds to truth values) of these models are taken too seriously, the theory will misfire precisely where it counts.

As we will see in the next chapter, the use of precise and fixed meanings (or the illusion of it) leads to more problems than just puzzles about vagueness. The error also seems to be at the heart of a broad range of puzzles at the heart of the analytic philosophical tradition.

# 5

## Consequences for Analytic Philosophy

In previous chapters I've argued that the static picture of the lexicon is bad philosophy and bad science, but the picture isn't just wrong. Its adoption in analytic philosophy has led to a series of puzzles that have plagued analytic philosophy for the past hundred years. In this chapter I look at a handful of these problems and show how the static picture has led us astray, and how the dynamic picture of the lexicon helps us to see our way clear of the puzzles generated by the static picture.

### 5.1 Contextualism in Epistemology

The classical skeptic challenges our claims of knowledge (for example, my claim that I know I have hands) by asking how we know that we aren't being deceived by an evil demon, or in more recent incarnations, that we aren't brains in vats being deceived by an evil scientist. If you can't know you are not a brain in a vat, then how can you know you have hands?

Contextualists approach the theory of knowledge by studying the nature of knowledge *attributions*—in effect, by examining our linguistic practice of attributing knowledge to people in different contexts. According to a contextualist like DeRose (1995), these linguistic practices suggest that there isn't a single standard of knowledge, but rather there are higher and lower thresholds for knowledge depending upon our context. In a casual conversation in a bar our standards of knowledge might be very low. In a courtroom the standards will be

much higher. In an epistemology class they might be even higher. The relevant standard of knowledge is usually taken to depend upon the context of the person who is reporting the knowledge state.

Let's add some detail to this picture and take a look at it from the perspective of the dynamic lexicon. Drawing on Hawthorne (2004: ch. 2) we can articulate a working definition of 'contextualism' in which the following two components hold.

- C1. [context-sensitive semantic values]: A given sentence, say 'Chesner knows that he has feet' has different semantic values relative to different contexts of utterance, (and this at least in part due to contextual parameters connected to the verb 'know' itself...). In brief, the contextualist claims that the epistemic standards required for someone to count as meriting a positive knowledge ascription varies from ascriber to ascriber, with the result that one ascriber may truly say 'He knows that he will be in Syracuse,' referring to a given person at a given time, and a different ascriber may say 'He doesn't know that he will be in Syracuse,' speaking of the same person at the same time.
- C2. [ascriber calls the shots]: According to standard contextualist semantics, the ascriber calls the shots, so to speak: the standards of application for the verb 'know' are determined by the ascriber and not by the subject (unless the subject happens to be identical to the ascriber).

As I observed in Ludlow (2005b), (C2) leaves matters open in a way that is not benign. For example, can the ascriber stipulate the context of utterance, or is it out of the ascriber's hands altogether? Is it the context alone that does the work? In this case (C2) would be sharpened as in (C2\*)—this appears to be the sharpening that Lewis (1996) was working with.

- C2\*. According to contextualist semantics, the ascriber's context of utterance calls the shots, so to speak: the standards of application for the verb 'know' are determined by the context in which the ascription is made and not by the context

in which subject appears (unless the subject happens to be identical to the ascriber).

I say that this appears to be Lewis's view because if you are in a high-standards context (say an epistemology class with a philosophical skeptic like Peter Unger) there is nothing you can say or do or whistle to invoke the low stakes context.

Alternatively, if you believe in Meaning Control, you might think that, even if I am in Unger's epistemology class, I can invoke a lower stakes context by stipulating it— "well by the standards of Ludlow's class I know that P." If this is our view, then we sharpen up (C<sub>2</sub>) as (C<sub>2</sub><sup>\*\*</sup>).

C<sub>2</sub><sup>\*\*</sup>. According to contextualist semantics, the ascriber's context of utterance calls the shots, so to speak: the standards of application for the verb 'know' are either explicitly stated or are determined by the context in which the ascription is made and not by the context in which subject appears (unless the subject happens to be identical to the ascriber).

Contextualism offers the following answer to the skeptic: The person who denies we know we have hands is sometimes right, but only in contained artificial circumstances, and our knowledge claims are otherwise preserved. More importantly though, the proposal allows us to think about the nature of knowledge in each different setting (court of law, scientific laboratory, etc.) and thus offers a research program in which epistemology is intimately connected with the individual sciences and social institutions. There is no single static standard of knowledge, but a sliding scale along which there might be different standards of knowledge for different institutions and activities. The problem is that there are some powerful objections to contextualism. Let's start with Stanley (2005).

Consider examples (1) and (2).

- (1) By the standards of Unger's seminar, I don't know I have hands
- (2) By the standards of Chemistry, the stuff in the Chicago River isn't water

## 132 CONSEQUENCES FOR ANALYTIC PHILOSOPHY

Stanley (2005, ch. 3) argues that these are not so very different in kind—that there is really nothing special about the standards talk with epistemic predicates. Indeed, it is arguable that we get that sort of standards talk with every predicate (counting common nouns as predicates here). Stanley’s point is that by showing this similarity we undermine contextualism, but another way of taking the point would be as illustrating that contextualism in epistemology is just a special case of the dynamic lexicon—it is the dynamic lexicon applied to epistemic language.

Let me illustrate this idea. Another observation we could make is that it is very natural to add emphasis on the verb ‘knows’ and ‘water’ in these cases.

- (1a) By the standards of Unger’s seminar, I don’t **KNOW**  
I have hands
- (2a) By the standards of Chemistry, the stuff in the Chicago River  
isn’t **WATER**

Note also that we could accomplish the same thing by using scare quotes or using the “pair of fingers” gesture to indicate scare quotes when we are speaking:

- (1b) By the standards of Unger’s seminar, I don’t “know”  
I have hands
- (2b) By the standards of Chemistry, the stuff in the Hudson River  
isn’t “water”

What is going on here? I want to suggest that there is a kind of metalinguistic operation going on; we are swapping in a nonstandard usage, or a homophone from a different language.

The idea is that not only can we create microlanguages on the fly (as argued in previous chapters), but we can also borrow from other microlanguages we have participated in or are aware of, so long as we signal what we are up to. One way of signaling that this is what is going on is by using expressions like ‘by the standards of’ or ‘in so-and-so’s terminology’, or ‘if you ask a Chemist...’ or ‘in the *parlance* of the Ungerian’.

The phenomenon involves what linguists call “code-switching” — incorporating terms from another language. So, for example, if I utter (3), I am code-switching in at least two spots. First, I am using a bit of French when I deploy ‘parlance’ (either because of the shade of meaning it affords or because I’m a snob and a show-off). Second, I am code-switching with the word ‘know’.

(3) In the *parlance* of the Ungerian, I don’t know I have hands

Code-switching isn’t quite the same as borrowing something from one language into another. The dead giveaway is that pronunciation is usually normalized in the case of borrowing, but not code-switching. If I want to borrow the word for the Italian snack called ‘bruschetta’ I will say [broo-shetta], if I am code-switching I will say [broo-sket’-ta]. The claim here is that we are code-switching on homophones from different microlanguages—although if people in an Unger seminar on skepticism pronounced ‘knowledge’ as [newledge] we might be inclined to adopt their pronunciation, as in (3’).

(3’) In the *parlance* of the Ungerian, I don’t [new] I have hands

That’s the basic idea, but of course this requires that we revisit the idea of microlanguages.

The idea would be that the term ‘knowledge’ is fleshed out in different ways in different conversational contexts—that is, in different microlanguages. In Unger’s philosophy seminar, for example, ‘knowledge’ may be understood to entail extremely robust justificatory requirements. In my seminar, it might have explicitly weaker requirements, and in some contexts the justificatory requirements may remain undetermined.<sup>1</sup>

This, by the way, helps us to avoid the sort of conclusion that Lewis was driven to. On *my* story, knowledge isn’t elusive at all. It is true that there are ways of fleshing out the term ‘knowledge’ —for example, Unger’s—according to which none of our beliefs fall into the extensions of the verb, but if we have negotiated another meaning for ‘knowledge’

<sup>1</sup> This view, or something very much like it, is advocated by Stine (1999).

we can certainly talk about Unger's fleshing out of the basic term without taking on the corresponding standards of justification ourselves.

But which one of us is really talking about *knowledge*? We both are of course! But when we say that both of us are talking about the same thing we are winding back to a more underdetermined neutral modulation of the term 'knowledge'. We can certainly argue about admissible sharpenings and correct modulations and someone can dig in their heels and say that only Unger's sharpening/modulation is admissible, but this is just semantic chauvinism, not so different in kind from someone who insists without argument that the only admissible sharpenings/modulations of 'planet' must include Pluto or that the only admissible modulation of 'marriage' must omit same-sex couples.

Stanley is aware of this move (he calls it "contextualism on the cheap") and objects to it. First of all, let's see why it is cheap contextualism. It is "cheap" by Stanley's lights because the shift in standards does not trade on a context-sensitive argument position for standards of knowledge, but rather because the meaning of the lexical items shifts from context to context. The view is cheap in the sense that there is nothing special about the context sensitivity of knowledge claims—they are context-sensitive in the same way that all of our lexical items are. The dynamic lexicon applies to all words, not just our epistemic vocabulary.

Stanley offers three objections to cheap contextualism. First, he argues that if cheap contextualism takes the form of a simple lexical ambiguity, we want to know why the ambiguity is not reflected by distinct words in some other languages. We can call this argument "The ambiguity manifestation argument". The idea (from Kripke 1977) is that canonically ambiguous words like 'bank' typically have their distinct meanings translated as different words in other languages. We should expect the same from any ambiguity we are prepared to posit. That is, a genuine ambiguity should be manifest as distinct words in some language of the world.

I don't think this is an argument so much as guiding methodological principle, but for all that it is a pretty good methodological principle and I'm happy to endorse it here. I don't think the proposal I'm making

here violates it, however, since I'm not suggesting a traditional ambiguity thesis—one in which the two meanings associated with the expression are radically different (and probably come from the accidental phonological convergence of two distinct words). It would be absurd to think that every modulation of a word should have a distinct morphological realization in some other standard language of the world.

Stanley's next two objections take the form of a dilemma. Once you set aside the ambiguity thesis, then you appear to be endorsing a "loose talk" version of contextualism. But loose talk how? Either you are thinking that there is a literal meaning and the loose talk diverges from that meaning but is licenced pragmatically as in Laserson (1999), or you are opting for some sort of semantical notion of loose talk in which the meaning is a semantically "roughed up" version of the precise meaning. This might involve a kind of "rough-it-up operator" in the sense of Barnett (2009).

Stanley (2005: 82) rightly points out that the first horn of the dilemma—saying that knowledge ascriptions are literally false but we are still entitled to assert them because they are pragmatically useful—effectively gives up the game. It is precisely the position that many skeptics take: Knowledge claims are literally false but assertable.

If we avoid the first horn of the dilemma and take the second, then Stanley is ready to drop the hammer on us, arguing that semantic "loose talk" becomes nothing more than deviation from the literal meaning (like saying 'It's 3 o'clock' when it is really 3:03) and if that's the case, we still collapse into something like a pretense theory or (again) a warranted assertability theory. The idea is that a roughly speaking operator is a kind of pretense operator in sheep's clothing.

Stanley's argument rests on a widely held assumption that I have rejected throughout this book. The assumption is that any notion of semantic loose talk is talk that is somehow parasitic on or generated from a more precise "correct" meaning (perhaps with a loosey-goosey operator which means something like "this rounds off to the truth"). So for example, '3:00 o'clock' really means 3:00 o'clock on the nose, but we understand that for certain purposes 3:03 is close enough, we apply the loosey-goosey operator and what we say is true. Alternatively, we

might think that the meaning of the term itself could be loosened up. For example, we could say that ‘flat’ means absolutely flat, but for certain purposes the meaning of the term can be relaxed to allow irregular surfaces within the extension. The problem for Stanley is that the theory of the dynamic lexicon does not make use of the basic assumption at work in his argument.

The dynamic lexicon theory diverges from such a ‘loose talk’ story in important ways. First, following the doctrine of Meaning Egalitarianism (introduced in section 3.2), we have rejected the idea that the core meaning is the absolute meaning (e.g. absolutely flat, or perfectly hexagonal, or known with Cartesian certainty)—there is no privileged core meaning. For example, the predicate ‘flat’ pertains to some underdetermined range of surfaces, ‘hexagonal’ to an underdetermined range of shapes, and ‘knows’ to an underdetermined range of psychological states. Technical uses of expressions like ‘absolutely flat’, ‘perfectly hexagonal’, and ‘Cartesian knowledge’, are just that—technical uses. They are modulations that have been introduced from time to time, but they are not the starting place for the semantics of ‘flat’, ‘hexagonal’, and ‘knowledge’.

As I noted in earlier chapters, this point even extends to cases like saying it is 3:00 o’clock. The expression ‘3:00 o’clock’ doesn’t mean precisely 3:00 o’clock on the nose unless we sharpen it to mean that for certain purposes (or try to do so—this may not be possible). The meaning that it is 3:00 on the nose is just one sharpening of the expression and not the core meaning. This point even extends to expressions like ‘now’. We needn’t take ‘now’ to mean a time slice at some exact time of utterance, it can just as easily include vast stretches of time (as in the ‘Universe is expanding now’). The “looser” meanings of ‘now’ are not derived from the precise meanings, but they are co-equal sharpenings of the meaning of the expression. In other words, there is nothing specious about “the specious present.”

Of course, as we saw in Chapter 4, there is a story about why we expect core meanings to be precise. Many theorists (e.g. Braun and Sider) make an unsupported assumption—that there is a kind of extreme precision in the meaning-stating vocabulary of semantic theories. For example, if we think of the semantic theory as a T-theory,

then the expectation is for the right-hand sides of the axioms and theorems to deliver precise meanings. And the expectation is not only for the theory to deliver meanings that are precise, but meanings that are simple or absolute in some sense. But as we also saw in Chapter 4, this expectation is mistaken. On the dynamic lexicon approach, if we provide a semantic theory for a microlanguage, then the meanings of the expressions on the right-hand sides of our axioms and theorems will be underdetermined. It is no virtue to introduce precision if it serves to mischaracterize the meanings of the expressions we are investigating. A good semantic theory must lift the underdetermination of the object language into the metalanguage of the theory (just as tensors lift tense into the metalanguage<sup>2</sup>). Adding precision to axioms and theorems is simply to misstate the general meanings of the expressions, which are not sharp, although they will be sharpened up (or loosened) to the degree necessary in various discourse contexts.

It should also be noted that the move to the dynamic lexicon requires that we sharpen up the definition of contextualism that we introduced in the beginning of this section. If we adjust the definition to reflect our understanding of the dynamic nature of the lexicon the result is the following (where changes are underscored).

- C1'. [context-sensitive semantic values]: A given sentence form, say 'Chesner knows that he has feet' has different sharpenings and modulations and hence different semantic values relative to different contexts of utterance. In brief, the contextualist claims that the epistemic standards required for someone to count as meriting a positive knowledge ascription varies from ascriber to ascriber, with the result that one ascriber may truly utter the form 'He knows that he will be in Syracuse', referring to a given person at a given time, and a different ascriber may truly utter the form 'He doesn't know that he will be in Syracuse', speaking of the same person at the same time.

<sup>2</sup> In other words, tensors believe that one has to use tensed language to give a characterization of the meaning of a tensed expression. There is no way to do this in untensed language.

## 138 CONSEQUENCES FOR ANALYTIC PHILOSOPHY

C2<sup>\*\*'</sup>. According to cheap contextualist semantics, the ascriber's context of utterance calls the shots, so to speak: sharpenings and modulations and hence the standards of application for the verb form 'know' are either explicitly stated or are determined by the context in which the ascription is made and not by the context in which subject appears (unless the subject happens to be identical to the ascriber).

The changes are necessary because it is open to dispute as to whether we are talking about the same lexical items (under different sharpenings) or distinct lexical items. The talk of forms is a way of staying neutral.

Once the definition of contextualism has been sharpened up to accommodate the dynamic lexicon, some of the additional objections to contextualism start to unravel. For example, Hawthorne (2004: s. 2.7) claims that there is a problem for any contextualist account constructed as above, and it turns on the following argument. Consider first, the following premises.

*The True Belief Schema (TBS)*

If S believes that *p*, then S's belief is true if and only if *p*.

*The False Belief Schema (FBS)*

If x believes that P, then x's belief is false if and only if it is not the case that P.

*Disquotational Schema for 'Knows' (DSK)*

If an English speaker E sincerely utters a sentence *s* of the form 'A knows that *p*' and the sentence in the that-clause means that *p* and 'A' is a name or indexical that refers to *a*, then E believes of *a* that *a* knows that *p*, and expresses that belief by *s*.

*True Belief Principle (TBP)*

If a speaker sincerely accepts an utterance *u* and *u* has a semantic value *p*, then the belief manifested by his sincerely accepting *u* is true iff *p* is true.

Hawthorne reasons as follows. Suppose a speaker S in a low standards or low stakes situation utters (4) to a hearer H in a high standards or high stakes situation:

(4) I know I have feet

Since H sincerely accepts and understands ‘I know that I have feet’ and since the semantic value of that utterance is true, TBP tells us that the belief H manifests by sincere acceptance is a true belief. Because H is a (standard) contextualist H is committed to saying that S’s belief is true. But the Disquotational Schema for Knows now forces H to say (5).

(5) You (S) believe that you know you have feet

but then, by the True Belief Schema H deduces (6).

(6) You (S) know you have feet

So, standard contextualism, in combination with TBP, TBS and the DSK would have H conclude that S knows S has feet. But this conclusion is supposed to be forbidden by the standard contextualist, because if H to sincerely accept ‘You (S) know you have feet’, then [H] would have a false belief since, in the scenario envisaged (where H is in a high stakes/standards situation), the semantic value of the latter sentence is false. The problem, Hawthorne alleges, is that I can only ever read (6) as, roughly

(6′) You know-by-my-standards you have feet

I don’t think Hawthorne’s argument touches the theory I have advocated, since when S and H create a microlanguage together they are supposed to be on the same page—that is, entrained—on the meanings of key terms. This would suggest that they ought to be synced up on what should count as ‘knowledge’. A low-stakes/high-stakes divergence should not—indeed cannot—happen within a single microlanguage. If one is detected the participants should pause and correct the situation. They can work with the low stakes version of ‘know’ or the

high stakes version, or they can lexically distinguish them somehow and deploy both.

Of course this only works if the meaning of ‘knows’ is modulable and in particular is modulable across a dimension involving our standards for knowledge, and Hawthorne believes this is simply not the case—he believes that ‘knows’ is a kind of strange outlier:

... we have few devices in ordinary life for implementing the clarification technique when it comes to ‘know’... We don’t have anything like the ‘of F’ and ‘for a G’ locutions available. Nor do we have anything like the hedge devices ‘roughly’ and ‘approximately’ available. (2004: 105)

I don’t think Hawthorne is right about this. We already talked about the case where I say ‘by the standards of Unger’s seminar’ and we routinely say things like ‘Do you know or do you *know*?’ when challenging someone’s knowledge claim. As I pointed out in Ludlow (2005b) a simple Google search will reveal that we use many constructions for hedging or qualifying knowledge claims. The following examples were the product of one hour with Google:

- known by any objective standards,
- known (by occidental science standards)
- known by earthly standards,
- knowing for sure what we "know" by academic standards
- know with some level of confidence
- know, with some reasonable certainty,
- know with some, albeit imperfect, reliability
- know with some degree of certainty
- know with some degree of accuracy
- know with some confidence
- know with some precision
- know with some authority.
- know with some probability
- know with some degree of authority
- know with some clarity
- know with some accuracy

- know with some level of confidence
- know with complete certainty
- know with complete certitude
- really know
- how well do we know
- pretty well know
- in effect knew
- sort of in effect knew
- by secular standards of knowledge.
- by today's standards of knowledge,
- by high standards of knowledge and accuracy.
- perceived standards of knowledge
- by your own standards of knowledge
- By today's standards of knowledge
- previously known standards of knowledge.
- with general contemporary standards of knowledge
- inappropriate standards of knowledge
- halachic standards of certainty
- legal and scientific standards of certainty,
- different standards of certainty are used in science and in politics.
- standards of certainty in law: criminal trials, civil lawsuits, government regulation, legislation.

It thus appears that the verb 'knows' (and epistemic adjectives like 'certain') has an extremely rich thematic structure which encodes not just the agent and the content of the belief, but potentially argument places for standards of justification and evidence, for subjective certainty of the report, for the reporter's responsibility for having and defending the knowledge, the source of the knowledge, and the mode of presentation of the content of the knowledge report. The resulting picture of this first approximation would be along the following lines.

*Know*

Agent: the ascriber,

Theme: the propositional content of the knowledge,

standards of justification: “by legal/scientific/etc. standards we know that that . . .”

standards of evidence: “by the evidential standards of criminal law we know that . . .”

degree of subjective certainty: “I know with confidence that . . .”

standards of subjective certainty: “by the standards of physics I know with confidence that . . .”

degree of responsibility for knowledge: “you know very well that . . .”

source of knowledge: [in many languages with “evidentials”], (“I see/know with my own eyes that . . .”)

mode of presentation: “John knows, in effect that p, although he wouldn’t agree to it in those terms

To put it another way, ‘knows’ is modifiable along many dimensions. Far from being a rigid outlier, it is one of our more flexible terms. Of course it isn’t modifiable in the same way that ‘flat’ is and cannot be modified with ‘for a’ and ‘roughly’, but that is presumably because it is a verb and not an adjective (notice how awkward it is to say ‘flatten it for a pool table’ and the meaning shift in ‘John roughly flattened the pool table’). Verbs show different marks of modification than do adjectives. There is of course the deverbal form ‘known’, but notice that deverbal adjectives are not modified in the way that adjectives like ‘flat’ are. Consider the awkwardness ‘it was flattened for a pool table’ and ‘it was roughly flattened’. I don’t think that is much better than ‘it was roughly known’.

Hawthorne has one additional argument, which trades on our ability to offer clarification when we are challenged in a knowledge claim. Hawthorne notes that we have three choices when one of our claims is challenged. We can stick to our guns, concede the point, or clarify. We offer clarification when there has been a misunderstanding. Here is one of Hawthorne’s examples.

*The glass is empty*

Challenge: Well it’s got some air in it

Reply: All I was claiming is that it was empty of Vodka

I agree that appeals to the meaning of ‘knows’ are rare when we are defending ourselves against being wrong. But surely this is because what we said is *false*, and if you believe that knowledge is factive, then it is sort of pointless to weasel out of your error by talking about the operative standards of justification. (In the example just given the glass *was* in fact empty of vodka.) At best you can diminish your culpability. But what about cases where we withhold knowledge?

Hawthorne seconds Feldman’s (2001: 77) claim that cases like this are rare, and I should think they would be. One problem is that for these cases to make any sense at all we need an instance where a knowledge claim is withheld in a high stakes situation, we come to know that the knowledge claim was withheld, and we confront the withholder in a lower stakes situation in which the knowledge claim would be true, oblivious to the change in stakes. The result is something like this.

*I know who the killer is*

Challenge: Yesterday you said you didn’t know who the killer was  
 Reply: Yah that was in court, and I didn’t know beyond a shadow of a doubt

*I know where the car is*

Challenge: Ha, an hour ago you said you didn’t know  
 Reply: Well I was considering brain-in-vat scenarios in epistemology class

The challenger needs to be someone obtuse or perhaps obnoxious, but surely no more obnoxious or obtuse than the person who insists the glass is not empty because it has air in it.

Summarizing this section, we can see that one of the leading puzzles in epistemology has the hallmarks of a classical case of meaning modulation. Once we recognize that modulation is taking place—that meaning is shifting between contexts (Unger’s classroom and the bar)—then we need to be careful about keeping track of which meaning is appropriate for the context we are in when we engage in knowledge attribution. There is no reason to deploy a meaning of ‘knows’ that is borrowed from a completely different context.

Ram Neta (2008) observes, correctly, that sauce for the goose is sauce for the gander and that if the dynamic lexicon is applicable to terms like ‘knowledge’ and ‘water’ then surely it is just as applicable to terms like ‘meaning’, ‘modulation’, ‘microlanguage’, etc.—in effect to all of the theoretical vocabulary introduced in this book. Does my position cannibalize itself? It is hard to see why; I said that underdetermination was lifted into the metalanguage, after all. The language of the metatheory is modulable just like everything else.

But perhaps the objection can be given sharper teeth: Someone—say Unger—might refuse to bite on the dynamic lexicon. He might argue in the following way: “when *I* talk about the meaning of a term I am using the term ‘meaning’ in a way that just doesn’t allow this sort of nonsense. The meaning of a word is what it is, I’m telling you about the meaning of ‘knows’, and it has an invariant meaning. This isn’t something that is up for grabs.”

One might anticipate a related version of this objection as well (this was first brought to my attention by Jennifer Lackey). It goes like this: Ludlow’s dynamic lexicon is fairly promiscuous—it looks like any sort of meaning goes. So while Ludlow might deploy a wimpy contextualist version of ‘knows’ there is no reason for Unger to do so. How on Earth are we going to move Unger (or any skeptic) off of his position? For that matter, why should we follow Ludlow on the meaning of ‘meaning’ or ‘modulation’ or ‘microlanguage’?

Both versions of the objection miss the mark. As we saw earlier, the mere fact that there is variation in the meaning of these expressions does not mean that anything goes and certainly does not mean that anyone is entitled to stubbornly dig in on the meaning of a term. To the contrary, the process by which we become entrained often involves argumentation, and argumentation is a normative activity. That is, norms govern the way we negotiate (litigate) word meanings. Someone who digs in on the meaning of ‘meaning’ is in no better position than someone who digs in on the meaning of ‘planet’ or ‘marriage’. They are appealing to a fact of the matter that simply does not obtain.

Now the Ungerian might reply, “who is digging in here? You dug in on a modulation of meaning that allows it to be dynamic.” But of

course I haven't dug in without argument; the static conception is a possible modulation of 'meaning' but it just isn't appropriate for a broad class of uses of 'meaning' by linguists and philosophers—or so I have argued in this book.

Another objection due to Jennifer Lackey (pers. comm.) now becomes salient. Given that we are in a position of arguing about the right modulation of 'knowledge' how have we advanced our position over the traditional debates in epistemology where contextualists and invariantists argued about the *concept* of knowledge? What's the difference? There is a two-part answer to this question.

First, we are often told things to the effect that contextualism is not really a theory about knowledge—knowledge itself—but is merely about how the *word* 'knowledge' is *used* in certain contexts. But this isn't right. After we modulate the meaning of 'planet' and begin deploying that term in our microlanguages we are talking about planets. After we modulate the meaning of the term 'persons' and deploy that term we are talking about persons. After we modulate the term 'knowledge' and deploy it, we are talking about knowledge.

It is certainly true that we aren't getting at planets, persons, or knowledge by first probing some abstract concept of planet, person, or knowledge. There is no mysterious concept of knowledge in Plato's heaven to which we as philosophers have special access.

The second part of the answer is sociological. Once we puncture the myth of the concept of knowledge and understand that we are in a conflict over the best way to modulate the meaning of the term 'knowledge' relative to some shared tasks or interests, we are more apt to be critical of arguments that rest heavily on the weight of authority. Philosophers since Descartes may have used the term 'knowledge' in an invariantist way, but why should we? What special claim do these experts have on us, so that we must feel compelled to reject the usual contextually sensitive uses of knowledge when we are in philosophical discussion? There may be a Cartesian tradition about the term 'knowledge' which takes it to have an invariantist meaning, but that tradition counts for little—or at least we can rightly ask why it counts for anything in our current discussions in epistemology.

This doesn't mean that the dynamic lexicon by itself leads directly to contextualism about knowledge, but it does three things for us. First, it provides us a plausible version of "cheap contextualism" that escapes the difficulties that have afflicted other versions of contextualism. Second, it neutralizes the claim that contextualism is defective because it only speaks to linguistic usage. Finally, it puts us in a position to challenge the semantic reach of philosophical tradition.

## 5.2 Paderewski, Peter, and Pierre

In the previous section we saw a case in which a term—'knowledge'—shifted meaning across contexts. We also saw that if we failed to notice this shift in meaning we were led into a philosophical puzzle. I believe a similar state of affairs holds in Kripke's (1979) case of the name 'Paderewski'.

Recall that on Kripke's story, someone, let's call him Peter, might come to believe that there are two Paderewskis—one a famous pianist, and the other a famous Polish statesman. Unknown to Peter, Paderewski the pianist just *is* the famous Polish statesman. Under such circumstances, Kripke asks, is it not possible that we could truly report Peter's beliefs both by an utterance of (7) and an utterance of (8)?

- (7) Peter believes Paderewski is a pianist
- (8) Peter believes Paderewski is not a pianist

And if that's the case, and if we can conjoin the that-clauses into a single belief, then we end up reporting that Peter inconsistently believes that Paderewski both is and is not a pianist.

Once we shift our perspective on the nature of the lexicon I think there is a natural solution to this puzzle. There are numerous ways to incorporate this insight, but I'm partial to a proposal in Larson and Ludlow (1993), which incorporates the dynamic approach to the lexicon into their account of attitude attributions.

On the Larson and Ludlow proposal the puzzle was to explain our ability to know when utterances of two distinct that-clauses attribute the same belief or count as saying the same thing, and alternatively

when the same that-clause expresses two distinct belief states (in different contexts). The first challenge was to show how using *different words* in a that-clause presented to different hearers (or the same hearer in two different contexts) could count as having attributed the *same attitude* to an agent. The second challenge was to show how using the *same words* with different hearers (or the same hearer in different contexts) could count as having attributed *different attitudes* to the same agent. Distilled to its essence: How do we use different words to express a single attitude, and how do we use the same words to express different attitudes?

Our proposal was offered in support of a particular theory of the semantics of belief, but it can be put in a theory-neutral way. The idea is this: When people report beliefs, they are not so much interested in trying to express something that is isomorphic to a representation in some agent's head; rather a speaker S and hearer H are collaborating together on constructing a way to talk about the mental life of some agent A. There are two basic elements to this; first, construct the theory of the agent's mental life, and second, construct a microlanguage in which to talk about this theory.

Sometimes, when reporting on beliefs we are indifferent to how an agent represents something; we only want the facts. If we are interested in whether Padrewski is a good piano player and Peter (who we trust) says that he is, then we are indifferent to how Peter represents Paderewski. Alternatively, if we are interested in the behavior of Peter (for example, whether he will go to a concert when he hears that the performer "will be the great statesman Padrewski") then how he represents things does matter. On the Larson and Ludlow proposal the idea was that each time S and H engage in a belief report they construct a microlanguage specifically for reporting on the beliefs of the agent of interest. They will tacitly entrain on the expressions to use in the characterization of the agent's cognitive states—or rather those states of interest to S and H.

What happens if we take into account the idea that words are introduced and word meanings are modulated and litigated on the fly by discourse participants when ascribing attitudes? Words are typically

introduced to be just fine-grained enough to resolve potential misunderstandings and ambiguities. For example, if two experimental participants are given a stack of pictures and are tasked with finding a joint way to refer to them, a stack of pictures with just one car is probably only going to yield the expression ‘car’ as the identifying expression for the picture. If there are multiple pictures, then the identifying expressions become more fine-grained, for example including qualifiers like ‘red car’ or ‘sports car’ or something even more fine-grained if necessary. Extending this result to the case of Peter, the prediction would be that a speaker S and hearer H would ordinarily employ distinct expressions for speaking of Peter’s distinct belief states just in case the two states are being discussed in a single conversation—that is, just in case the theory of Peter’s beliefs relative to our shared interests calls for us to introduce a temporary lexical distinction.<sup>3</sup>

The Kripkean assumes that there are contexts in which it might be said both that “Peter believes Paderewski is a piano player” and “Peter does not believe that Paderewski is a piano player.” But are there really such contexts? We know there are contexts in which the former attribution might be made, and we know there are contexts in which the latter attribution might be made, but are there natural contexts in which *both* reports might be made? Even if speaker and hearer are aware of the agent having multiple lexical entries for Paderewski, there is no way for the speaker to communicate separate beliefs in a single context without somehow signaling which entry is under discussion at the moment. The experimental evidence cited earlier (e.g. Wilkes-Gibbs and Clark 1992; Metzinger and Brennan 2003) suggests that in such cases speakers will refine the expressions used, perhaps as ‘Paderewski qua pianist’ and ‘Paderewski qua neighbor’ or by some other mechanism.

If this is right, then there is a fallacy at work in Kripke’s “Puzzle about Belief.” The fallacy involves the conjunction of two sentences that have the appearance of contradicting each other (they have the form Fa and  $\sim$ Fa) but they do not contradict because they come from different

<sup>3</sup> Other experiments involve the naming of tangrams. See Wilkes-Gibbs and Clark (1992), Metzinger and Brennan (2003).

microlanguages. The fallacy, which in formal logic is obvious, is made here because we think of English as a kind of static object that we are all speaking. If I'm right, it isn't an external object that we learn to speak and it isn't static at all—the lexicon can be bent to form slightly different “languages” in different contexts (and for different discourse participants in those contexts). As we saw in Chapter 4, logic (when applied to natural language) needs to be sensitive to word meaning modulation.

A similar response is also available for the Kripke London/Londres case. In this case an agent, Pierre, is raised hearing about how beautiful London is (perhaps he has also seen pictures of beautiful London scenery). He comes to have a belief that he expresses with the phrase ‘Londres est Jolie’, and we translate his words and report him as believing that London is pretty. But then he is kidnapped and taken to a sketchy part of London that is quite grim. He learns that his new city is called ‘London’ and does not realize it is the place he had heard about in Paris and that he had been calling ‘Londres’. He now says ‘London is not pretty’ and we report him as believing that London is not pretty.

The problem is that Kripke's way of posing the dilemma oversimplifies the situation from the perspective of the dynamic lexicon. What, for example, do translators do?

They surely don't translate from one fixed “source language” to another fixed “target language”. In Chapter 3 we saw why this was a failure in an attempt to translate into a computer language. To see why this is also a bad idea in the translation of natural languages consider the situation faced by two Serbian friends of mine who are translators working on translations from English into Serbian. One was translating Tolkein; the other was translating *The Color Purple*. So, exactly how does one translate Elvish expressions or Rural Black English Vernacular into Serbian? One common and very unhappy strategy in Serbia is to translate Black English Vernacular into Bosnian; I think we can agree that this is not the right answer. In point of fact, translators are not in the business of translating from source to target so much as extending and morphing the target language so as to communicate the

## 150 CONSEQUENCES FOR ANALYTIC PHILOSOPHY

ideas found in the source. Pannwitz (1917) had an interesting insight on this score:

The fundamental error of the translator is that he stabilizes the state in which his own language happens to find itself instead of allowing his language to be powerfully jolted by the foreign language. (From Venuti 1995: 148)

Of course on my view it is not that the translator's language is changing so much as the translator is establishing a microlanguage with the readers of the so-called translation. Direct coordination is out of the question with written translations, but assumptions about the knowledge and background of the audience will direct the way in which the microlanguage is constructed. In the case of a translation for a co-present audience member there should be no issue; if it is relevant to our modeling of Pierre's mental life, we need to introduce distinct names for London into the microlanguage.<sup>4</sup>

In the case of Pierre we have a genuine microlanguage in the works; we want to communicate something about Pierre's mental state and we work with our communicative partners to do just that. We select terms that will help the discourse participants to construct the relevant theory of Pierre's mental state. Sometimes, we will use the name 'London' to describe Londres-beliefs (the beliefs he expresses using the term 'Londres') and sometimes we won't. That is OK, because there is no fixed meaning to 'London.' In some microlanguages 'London is pretty' expresses his beliefs and in other microlanguages 'London is pretty' does not express his beliefs.

Again I think that this is a case where Kripke's argument rests on a kind of equivocation. If the meaning of an expression like 'London' was indeed a fixed and fully determined common coin in a broadly shared abstract language, then Kripke's question would certainly make sense, but if we think in terms of lots of little microlanguages the question doesn't really make sense. It is like taking a symbol like 'v' and, noting that it has a different truth table in different logics, asking is it or is it not to be understood as an inclusive classical disjunction. Well, in some logics it is and others it isn't.

<sup>4</sup> See Richard (1990) for a similar point about translation.

At one point Kripke frames up the question this way: “Does Pierre, or does he not, believe that London (not the city satisfying such-and-such descriptions, but *London*) is pretty?” In doing this, with the stress on ‘London’, there are a couple of points Kripke could be trying to get across. First, by stressing ‘London’, he could be asking does he not believe that *it* is beautiful. And here of course the answer is “both.” With respect to the object itself, Pierre has inconsistent beliefs. There is no puzzle in this. When we reintroduce the sense component we can see why Pierre’s beliefs are rational, and why he behaves the way he does and says the things he does.<sup>5</sup>

Alternatively, he could be asking “should I or should I not characterize Pierre’s belief with the expression ‘London is pretty?’” In this case the answer is: It depends.

Notice that I introduced the term ‘sense’. I believe that words can express senses, and that they express different senses on different occasions of use (sense and shades of meaning are precisely the stuff that is up for play in the negotiated construction of microlanguages). But if this is right we see that a series of puzzles about the nature of indexicals come to lose their bite.

### 5.3 Understanding Indexicals

In the previous two sections we saw cases (the words ‘know’ and ‘Paderewski’) where occurrences of terms across different contexts were falsely assumed to have the same meaning across those diverse contexts, and we saw how those assumptions led to philosophical puzzles. It turns out we can also identify cases in which terms having different forms are *falsely* assumed to have different senses (and hence different meanings).<sup>6</sup> As we will see, even if two terms routinely are used to express different senses there can be contexts (and microlanguages) in which the terms can be used to express the very same sense.

<sup>5</sup> What happens to these senses when the sentence is evaluated in other possible worlds? Presumably one needs a theory of modal discards here. In other words, counterfactual evaluation calls for modulations that thin out the sense content of the word.

<sup>6</sup> Notice that here I am understanding ‘meaning’ to be inclusive of sense content.

Let's begin by considering John Perry's (1977) famous objections to Frege on demonstratives. The point of departure for Perry's discussion is the following passage from Frege (1956: 296).

If someone wants to say the same today as he expressed yesterday using the word 'today', he must replace this word by 'yesterday'. Although the thought is the same, the verbal expression must be different so that the sense, which would otherwise be affected by the differing times of utterance, is readjusted. The case is the same with words like 'here' and 'there'. In all such cases the mere wording, as it is given in writing, is not the complete expression of the thought, but the knowledge of certain accompanying conditions of utterance, which are used as means of expressing the thought, are needed for its correct apprehension. The pointing of fingers, hand movements, glances may belong here too. The same utterance containing the word 'I' will express different thoughts in the mouths of different men, of which some may be true, others false.

Perry argued that Frege gets into trouble by trying to identify the sense of a sentence (utterance) with a thought. Why? Well, because 'yesterday' and 'today' presumably have different senses, and it therefore follows that 'Today is a fine day' and 'Yesterday is a fine day' must have different senses (since they are *composed* of different senses). But if I can express the same thought today with an utterance of 'yesterday is a fine day' that I expressed yesterday with an utterance of 'today is a fine day' then thoughts cannot be associated with senses. Different senses are deployed in expressing the same thought so thoughts are not in a one-to-one correspondence with the senses of sentences.

It seems that Frege has to give something up. To help us get clear on what the options are, Heck (2002) lays out the problem space as follows, suggesting that Frege was committed to the following doctrines.

- (1) There can be different Thoughts that "concern the same object" and ascribe the same property to it. For example, the Thought that Superman flies and the Thought that Clark Kent flies are different, even though Superman is Clark Kent.
- (2) Sentences of the form 'N believes that *a* is *F*' and 'N believes that *b* is *F*' can have different truth-values, even if '*a*' and '*b*' refer to the same object.

- (3) Sense determines reference.<sup>7</sup>
- (4) The sense of a sentence is what one grasps in understanding it.
- (5) The sense of a sentence is a Thought.

One way out is of course to reject 2; referentialists hold that the truth values are literally the same, but the sentences are put to different uses. Heck's way out was to reject (4)—that is, to reject the idea that there is a single thought associated with the understanding of a sentential utterance. But before we opt for these or other ways out of Frege's problem, it is important to see that another assumption is required to generate the problem. That is, there is a sixth premise, which we can put the following way.

- (6) A sense is intimately tied to its form of expression, so that the senses of 'today' and 'tomorrow' remain constant.

As we saw in the case of 'Paderewski', on the dynamic lexicon view, the sense of a term can shift across contexts—in one context it may express the great statesman sense and in another context it may express the great piano player sense. Accordingly this premise needs to be regarded with some suspicion.

Everyone agrees that the *referent* of 'yesterday' shifts—every day it picks out a new day after all. However, I think the standard view is that the *sense* of yesterday remains constant. The sense might be something on the order of "the day before the day of the utterance" (other options are of course available). However I don't believe this is the complete story of how demonstratives work; I think that not only does the referent of 'yesterday' shift, but the sense shifts as well. It shifts in that we "recalibrate" 'yesterday' so as to pick out the previous day's perspectival position. Specifically, the sense of yesterday is recalibrated every day to express a new perspectival position.

Here we have a different kind of modulation at work. We could call it "automatic modulation" because we don't need to litigate the new

<sup>7</sup> Heck (2002: 4) allows that this may be understood in a weak way: "On the weakest interpretation of (3), it speaks of 'determination' only in a mathematical sense: it claims only that senses are related many-one to references."

meaning at each step. The shift happens automatically if we are sensitive to our environment and the changes taking place in it, and in particular sensitive to changes in our perspectival position within the environment.

How does this work? Branquinho (2006) has suggested that, in cases like this, we have a kind of natural realignment, so that as we move from day  $d_1$  to day  $d_2$ , we shift the sense of the expressions ‘today’ and ‘yesterday’—that is, they do not merely refer to different days at different times, but they display different senses at different times. Here is how Branquinho puts the point.

Cases where one is dealing with indexical contents are problematic because they often involve some realignment in the linguistic means of expression of a thought—on the part of a given thinker—as time goes by. In other words, there are situations in which the verbal expression of an indexical thought entertained by a thinker at a given time must, at a later time, be readjusted in a certain way by the thinker in order for the thought in question to be then entertained; so that one could presumably say that some attitude held at the earlier time towards the thought in question has been retained by the thinker at the later time, the very same thought being the object of the attitude on both occasions. Naturally, such readjustments are to be thought of as being operated in the linguistic means employed for the expression of the thoughts. It does not make much sense—at least in the light of the picture of content we are assuming—to think of the thoughts as being themselves subjected to any sort of change or realignment.

Here is a way to think about the proposal. To express an egocentric thought had at time  $t_1$  at a later time  $t_2$ , we need to refer to that thought from our new egocentric position. So if there is a thought I express on day  $d_1$  as “Today is a fine day” the way I express *that very same thought* on the following day is to say “Yesterday was a fine day”. We keep track of the thought in our egocentric space. As we move through time, we shift the sense of ‘yesterday’ so that it not merely picks out the same day that ‘today’ did on the day before, but it orients our perspective in time so that it also locates the thought that we had the day before egocentrically—albeit from our new egocentric perspective. To do this, the sense of ‘today’ must be rebuilt every day so as to pick out the new egocentric perspective of that day, and ‘yesterday’ must be rebuilt every day, so as to pick out the previous day from a new egocentric

perspectival position. Clearly this Fregean notion of sense is a much thicker notion of sense than Kaplan's (1977, 1979, 1990) notion of character (which is stable across contexts and not perspectival in the sense I am interested in). As I said, this shift in sense across contexts isn't the product of negotiation, and it clearly must be largely automatic, but it is a recalibration for all that. It is a shift in the sense content of the meaning of the term.

Of course there are cases where this automatic recalibration can misfire, and it has been suggested that the possibility of these misfirings would make it impossible for us to express some thoughts that intuitively are quite salient to us.

Consider Kaplan's case of Rip van Winkle, who goes to sleep one day saying to himself, "today was a fine day." When he awakens twenty years later, he may want to again express the thought that he first expressed by the utterance that he made just before he fell asleep twenty years earlier. He may try to express this thought by saying 'yesterday was a fine day', but in doing so Rip *fails* to express what he did with his original utterance because he has lost track of the relative temporal position of his original utterance. To put it in terms of the proposal I offered, Rip thinks that the automatic recalibration of the sense of 'yesterday' is sufficient for him to allow him to express the perspectival thought he entertained when he was about to fall asleep. But of course it isn't sufficient. It misses the mark by a wide margin. So he seems to be unable to express the thought that he had. Evans thought Rip even lost the original belief:

I see no more strangeness in the idea that a man who loses track of time cannot retain beliefs than in the idea that a man who loses track of an object cannot retain the beliefs about it with which he began. (1981: 87n.–88n.)

But do we need to bite this bullet? It can be conceded that Rip has not successfully expressed the thought he had twenty years before (assuming he is shifting the meaning of 'yesterday' in the usual way), but it need not follow that he fails to retain or even lacks the resources required to express the thought he had previously.

To see the way out, let's begin with Branquino's example of the less extreme case where Jones (our modern-day Rip) utters (9) just before

midnight and, not realizing that midnight has just passed does not assent to (10).

(9) Today is a fine day

(10) Yesterday was a fine day

... if Jones mistracked time in the way described before, then ... what he would not be in a position to do at 00:01 a.m. on  $d+1$  is to T-retain the particular belief he held at 11:58 p.m. on  $d$  when he accepted [(9)]. In other words, he would not be able to re-express then such a belief by using a temporal indexical such as 'yesterday'. But it does not follow that it would be impossible for him to retain, or even re-express, *tout court* that belief..

Although Jones is *ex hypothesi* unable on  $d+1$  to keep track of  $d$ , i.e. to think of  $d$  as *yesterday*, he might still be said to have retained on  $d+1$  his former belief about  $d$  in a certain way, namely by means of memory, and not in virtue of the particular position he occupies in time or of his knowledge of such a position. One should therefore regard as unacceptable the claim that a disposition to accept [(10)] on  $d+1$  is necessary for a thinker to retain or re-express then a belief she had on  $d$  by accepting [(9)].

... a way by means of which a thinker like Jones could re-express on  $d+1$  the belief held on  $d$  (by accepting (12) then) would be to accept, or to have a disposition to accept, on  $d+1$  a token of a sentence such as

[(11)] That day was fine.

The demonstrative phrase 'that day' would be here taken as expressing in the context a memory-based demonstrative mode of presentation of  $d$ , i.e. a way of thinking of a certain day anchored upon a memory demonstration of the day in question.

Branquinho's point, I take it, is that there is a difference between an accidental failure to express the thought had earlier and a permanent inability to do so. Once informed of a passage of time (either a day or twenty years) we clearly regain the ability to express the thought had earlier. And even before we are informed of the time passage it is not clear that we lack the ability to express the earlier thought. For example, the sense of 'that day' can be appropriately tasked to express the sense of 'today' on the original tokening of the thought. Branquinho thinks 'that day' picks out a memory-based thought, but I think that there are other options. For example, we could take 'that day' to have a sense akin to "that day an unspecified number of days prior to my current temporal perspectival position."

The dynamic lexicon affords us the resources to use common expressions (like ‘that thing’ or ‘that day’) with new senses in addition to new referents. The senses that we attach to these expressions can depend not just on our current spatio-temporal perspective, but we can also adjust and modulate the senses to allow us to express memories from earlier spatio-temporal perspectives. Some expressions, like ‘yesterday’, are not modulable enough to always successfully do this (e.g. in Rip van Winkle cases) but other expressions (like ‘that day’) are sufficiently modulable and can be pressed into service in new micro-languages to express thoughts that were previously entertained, even though we have lost track of our relative spatio-temporal position. Of course, Rip needs us to play along and allow him to attach the relevant sense to ‘that day’, but there is no barrier to our doing so.

All of this is related to a point I made in Chapter 3. While we use language to express thoughts, and perhaps language is even prior to thought itself, it does not follow that we are imprisoned by our language and that it limits what we can and cannot think. Language is dynamic, and the resources we are afforded for expressing new *and* old thoughts are robust. Language is not a prison that prevents us from thinking new things, nor does it blockade us from thinking perspectival thoughts once entertained in distant and very different contexts. To the contrary, it seems to be a resource that is particularly well equipped to allow us to do all of these things. In fact, as we will see in the next chapter, it can also do quite a lot more.

## 6

## Metaphor and Beyond

A number of theorists (e.g. Atlas 2005; Carston 1997, 2002; Sperber and Wilson 1998; Wilson 2003) have challenged the distinction between the figurative and literal, in some cases assimilating the figurative uses to robust application of the mechanisms of word meaning modulation. So, for example, we might take a metaphor as being a very wide modulation of a word or phrase (or passage). I think this view makes good sense, and in this chapter I try to develop the idea within the framework of the dynamic lexicon.

It is easy enough to see why some philosophers and linguists think that the literal/metaphorical distinction doesn't hold up. If meanings can be modulated then surely there are contexts in which meanings can be extended in ways that we ordinarily classify as metaphorical.

Consider the following examples.

Love is a burnt match skating in a urinal.

(Hart Crane)

Love is something that hangs up behind the bathroom door and smells of Lysol.

(Hemingway)

There is no reason why we can't broaden the range of 'burnt match skating in a urinal' or 'something that hangs up behind the bathroom door and smells of Lysol' to be true of the human emotion love. Of course this raises the question of how these modulations of meaning hold the metaphorical senses that they do. How does modulating 'burnt match skating in a urinal' so as to include love accomplish anything? Or for that matter, how do we make sense of the idea that metaphor can play

an important role in the growth and development of scientific theory, as Hesse (1966) has argued. Metaphors are keys to scientific progress. But how?

I believe that the answer to both of these questions lies in the dynamic lexicon and in the norms of word meaning litigation discussed in Chapter 2. Those norms give us a way of deriving metaphorical meaning (and focusing scientific inquiry). The idea is this. When one is asked to modulate outside of the typical range of modulations (for example, modulating so that love counts as being in the range of ‘burnt match’) one is in effect being asked to suppose that the norms of word meaning litigation are being respected and, given this supposition, to deduce what salient and important properties are shared between the elements that are the subject of the modulation.

For example, as we saw in Chapter 2, the alternative modulations of the term ‘planet’ were justified by appeal to certain underlying physical properties that were taken to be important. The alternative modulations of the term ‘rape’ were justified by appeal to underlying social properties. When someone introduces a metaphor like ‘love is a burnt match’, we are being asked to modulate well out of the usual range, but it does not follow that the norms are being ignored. We are being invited to consider what important properties are shared between love and burnt matches skating in urinals (presumably in this case the property of being quite useless and more than a little bit disgusting). All of this is consistent with existing work on metaphor by Glucksberg (2001, 2003) and Glucksberg, McGlone, and Manfredi (1997).

There are also some similarities to Davidson’s (1978) proposal that metaphors mean what they literally mean but we are being asked to simply compare two things. The difference here is that the literal/metaphorical distinction has collapsed; the metaphorical meaning is simply a very wide modulation of word meaning. What makes that modulation useful is the recognition of shared properties.

This approach is particularly attractive when we think of it in terms of Hesse’s (1966) work on the role of metaphor in scientific theorizing. Very clearly the introduction of a metaphor can guide us in looking for and being attentive to shared properties. Thus, to use

one of Hesse's examples, a metaphor like 'the mind is a computer' encourages us to modulate the meaning of 'mind' and 'computer' so that the former falls in the range of the latter. If we go along with this metaphor we are being invited to seek out important shared properties, for example to seek out computational properties in human cognition (perhaps this leads us to hypothesize mental processes that are recursive or at least compositional, systematic, and that respect inferential coherence). Perhaps the metaphor becomes so productive that we even stop thinking about it as a metaphor. Or more accurately, perhaps the "metaphorical" modulation becomes accepted as a pretty typical modulation of the word.<sup>1</sup>

By the way, Hesse has remarked that all meanings are metaphorical, and while I think that claim is definitely wrong, it has an element of insight to it (did she intend her claim to be metaphorical?). If I (and other people exploring this idea) are correct, then Hesse is right that there isn't a big difference between the metaphorical and the literal. Every word is subject to meaning modification (a point that Hesse herself stressed). But it isn't so much that everything is metaphorical so much as there really is no interesting metaphor/literal divide. A metaphor is simply a word meaning modulation that reaches further than a typical modulation would.

The examples I've considered so far involve cases where we were working with metaphors that invoked predicates like 'is a burnt match', but what happens when the metaphor asks us to compare two individuals or objects—the classic example being 'Juliet is The Sun'. Following work by Burge (1973), Larson and Segal (1995), and more recent work

<sup>1</sup> When this happens—when the metaphor "dies"—we don't fossilize the invitation to compare categories of things, we rather (as noted by Stern 2000) see the effect of metaphors qua demonstratives. To illustrate this, consider a dead metaphor like 'mouth of the bottle'. The initial modulation of this phrase invites us to consider important underlying properties shared between the mouth of an animal and the opening in the bottle, but once the modulation is taken on as routine it becomes a predicate the range of which includes openings in typical bottles as canonical instances. It now merely denotes these objects. Is the invitation to compare part of the meaning of the metaphor or is it pragmatic? The dynamic lexicon approach is fundamentally neutral on this approach I believe.

by Gray (2012) and Fara (2012), we can take names to be predicates. What is it to modulate one of these predicates so that it includes another individual within its range?—for example, to modulate ‘The Sun’ so that Juliet falls within its range? In this case presumably we are asked to consider a predicate with a range that includes two members sharing a number of properties. We are then invited to consider what are the important shared properties given current interests (presumably the immediate interests of Romeo and the audience).

Should we worry that this process is wildly unconstrained? I don’t think so. We aren’t being asked to compare two things in any way we choose (as in Davidson 1978). We are being asked to seek out shared important properties on a particular dimension of interest. If the metaphor is applied in the physical realm (e.g. using the billiard ball model in describing the behavior of gases) then shared physical properties are in order. If the metaphor is applied in the social or personal realm then social properties are in order.

But the constraints are not exhausted here. In Chapter 3 I suggested that word meaning modulation may well be constrained by the thematic structure of verbs (agent, patient, theme, etc.). If this is right, then we can hypothesize that meaning modulation—even taken to the metaphorical level—will respect thematic relations.

We can illustrate this idea by slightly modifying an example from Kittay (1990), who discusses Plato’s use of the metaphor of Socrates as midwife. Kittay believes she is making the case for semantic fields, but it is pretty clear that the heavy lifting is being done by the thematic relations, which are preserved when the meaning of the predicate ‘is a midwife’ is modulated to include Socrates.<sup>2</sup> The thematic structure of that predicate would be as follows.

<sup>2</sup> If you find it problematic to think of the simple predicate ‘is a midwife’ as having this robust thematic structure, you can invent a verb which characterizes that which the midwife does—‘midwifing’ for example. Thus a midwife is someone who midwifes, and midwifes midwife states of affairs in which an agent brings about a result using an instrument. Following considerations in Larson (1998) I happen to believe that nouns like ‘midwife’ have very complex internal predicate structure that encodes these relations.



and other non-representational aspects of metaphor. Of course the same could be said for wide modulation more generally—I would imagine that the goal of modulation need not be purely representational; it could also be to express affect.

If this is right, then metaphor is simply an aggressive form of meaning modulation—designed to either make new concepts accessible to us or to allow us to refine the concepts already in play. I've also left the door open to the idea that modulation can help us to express certain kinds of attitudes that might otherwise not have been easily expressible for us on a standard modulation. But the power of meaning modulation to do this seems to undercut assumptions that have been made about the ways in which we are restricted by our language.

I think it is important to point out that even if others are in a power relation with respect to us we are never compelled to accept their modulations. We are never *compelled* to defer—we are never prisoners to our own or someone else's "language." On this point I take exception to the following passage from Deleuze and Guattari, cited in Venuti (1995: 273).

How many people today live in a language this is not their own? Or no longer, or not yet, even know their own and know poorly the major language that they are forced to serve? This is the problem of immigrants, and especially of their children, the problem of minorities, the problem of a minor literature, but also a problem for all of us: how to tear a minor literature away from its own language, allowing it to challenge the language and making it follow a sober revolutionary path? How to become a nomad and an immigrant and a gypsy in relation to one's own language.

We cannot be imprisoned by something that does not exist (that is by a fixed language). On the other hand we *do* need to be cautious in when and how we defer to the linguistic practices of our discourse partners, and we need to insist that semantic deference be paid only when it is warranted. More to the point, when necessary we should not hesitate to modulate word meanings aggressively so as to express the ideas (and feelings) we want to express.

This doesn't mean that anything goes. The proper response to attempts at regimenting our lexicon is not to retreat into word games, but rather to modulate meanings aggressively where

## 164 METAPHOR AND BEYOND

appropriate, but to at the same time exercise care and creativity in clarifying the modulations we wish to deploy, explaining their usages, and making sure that our discourse partners understand and respect these usages. Care and clarity do not hinder the creative use of language; they facilitate it.

## Bibliography

- Ajdukiewicz, K., 1974. *Pragmatic Logic*, tr. O. Wojtasiewicz (Dordrecht: D. Reidel and PWN—Polish Scientific Publishers; English tr. of *Logika pragmatyczna*, originally published 1965).
- Arana, A., and P. Mancosu, 2012. “On the Relationship between Plane and Solid Geometry,” *Review of Symbolic Logic*, 5: 294–353.
- Armstrong, J., and J. Stanley, 2011. “Singular Thoughts and Singular Propositions,” *Philosophical Studies*, 154: 225–54.
- Asher, N., 2010. *Lexical Meaning in Context: A Web of Words* (Cambridge: Cambridge University Press).
- Ashworth, E. J., 1974. *Language and Logic in the Post-Medieval Period* (Dordrecht: D. Reidel).
- Atlas, J. D. 2005. *Logic, Meaning, and Conversation* (Oxford: Oxford University Press).
- Augustine, R., 1991. “Marriage: The Safe Haven for Rapists,” *Journal of Family Law*, 29: 559–90.
- Bach, K., 1994. “Conversational Implicature,” *Mind and Language*, 9: 124–62.
- Baker, M., 1988. *Incorporation: A Theory of Grammatical Function Changing* (Chicago: University of Chicago Press).
- Barker, C., 2002. “The Dynamics of Vagueness,” *Linguistics and Philosophy*, 25: 1–36.
- Barnett, D., 2009. “Is Vagueness Sui Generis?” *Australasian Journal of Philosophy*, 87: 5–34.
- Barsalou, L., 1987. “The Instability of Graded Structure in Concepts,” in U. Neisser (ed.), *Concepts and Conceptual Development* (New York: Cambridge University Press), 101–40.
- 1993. “Flexibility, Structure, and Linguistic Vagary in Concepts: Manifestations of a Compositional System of Perceptual Symbols,” in A. Collins, S. Gathercole, A. Conway, and P. Morris (eds), *Theories of Memory* (Hove, UK: Lawrence Erlbaum Associates), 29–101.
- Beakley, B., and P. Ludlow (eds), 1992. *The Philosophy of Mind: Classical Problems/Contemporary Issues* (Cambridge, MA: MIT Press).
- Bezuidenhout, A., 2002. “Truth Conditional Pragmatics,” *Philosophical Perspectives*, 16: 105–34.
- 2006. “The Coherence of Contextualism,” *Mind and Language*, 21: 1–10.

## 166 BIBLIOGRAPHY

- Black, M., 1962. "Metaphor," in M. Black (ed.), *Models and Metaphors* (Ithaca, NY: Cornell University Press), 25–47.
- Blackmun, Justice H., 1973. *Opinion of the Court, Roe v. Wade*, 410 U.S. 113–78.
- Blackstone, W., [1765] 1859. *Commentaries on the laws of England* (Philadelphia: Lippincott).
- Bloom, P., 2000. *How Children Learn the Meanings of Words* (Cambridge, MA: MIT Press).
- Blutner, R. 1998. "Lexical Pragmatics," *Journal of Semantics*, 15: 115–62.
- Boguraev, B., and T. Briscoe (eds), 1989. *Computational Lexicography for Natural Language Processing* (London: Longman).
- Branquinho, João, 2006. "On the Persistence of Indexical Belief." Paper given to the APA Eastern Division Meeting, Dec. 2005, revised. Available online at <<http://pwp.netcabo.pt/0154943702/persistence.pdf>>.
- Braun, D., and T. Sider, 2007. "Vague, So Untrue," *Noûs*, 41: 133–56.
- Brennan, S. E., 1998. "The Grounding Problem in Conversation with and through Computers," in S. R. Fussell and R. J. Kreuz (eds), *Social and Cognitive Psychological Approaches to Interpersonal Communication* (Hillsdale, NJ: Lawrence Erlbaum), 201–25.
- and H. H. Clark, 1996. "Conceptual Pacts and Lexical Choice in Conversation," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22: 1482–93.
- and M. Schober, 1999. "Uhs and Interrupted Words: The Information Available to Listeners," in *Proceedings of the 14th International Congress of Phonetic Sciences, Satellite Meeting on Disfluency in Spontaneous Speech* (Berkeley, CA: 14th International Congress of Phonetic Sciences), 19–22.
- Brogaard, B., 2008. "Sea Battle Semantics," *Philosophical Quarterly*, 58: 326–35.
- Brownmiller, S. 1975. *Against our Will: Men, Women, and Rape* (New York: Simon).
- Budzynska, K., 2006. "Persuasion: The Practical Face of Logic," *Croatian Journal of Philosophy*, 4: 343–62.
- Burge, T., 1973. "Reference and Proper Names," *Journal of Philosophy*, 70: 425–39.
- 1979. "Individualism and the Mental," *Midwest Studies in Philosophy*, 4: 73–121.
- 2003. "Perceptual Entitlement," *Philosophy and Phenomenological Research*, 67: 503–48.
- Camp, E. 2006. "Metaphor and That Certain 'Je Ne Sais Quoi,'" *Philosophical Studies*, 129: 1–25.
- Cappelen, H., and E. Lepore, 2004. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism* (Oxford: Blackwell).

- Carston, R. 1997. "Enrichment and Loosening: Complementary Processes in Deriving the Proposition Expressed?" *Linguistische Berichte*, 8: 103–27.
- 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication* (Oxford: Blackwell).
- Chalmers, D., 2011. "Verbal Disputes," *Philosophical Review*, 120: 515–66.
- Chereminsky, E., 1982. "Rationalizing the Abortion Debate: Legal Rhetoric and Abortion Controversy," *Buffalo Law Review*, 31: 107–64.
- Chomsky, N., 1986. *Knowledge of Language* (New York: Praeger).
- 1994. "Naturalism and Dualism in the Study of Language and Mind," *International Journal of Philosophical Studies*, 2: 181–209.
- 1995. "Language and Nature," *Mind*, 104: 1–61.
- Clark, H., 1992. *Arenas of Language Use* (Chicago: CSLI Publications and University of Chicago Press).
- Cohen, S., 1991. "Skepticism, Relevance, and Relativity," in B. McLaughlin (ed.), *Dretske and his Critics* (Oxford: Blackwell Publishing), 17–37.
- 1999. "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives 12: Epistemology* (Oxford: Blackwell Publishing), 57–89.
- Condit, C., 1990. *Decoding Abortion Rhetoric: Communicating and Social Change* (Urbana, IL: University of Illinois).
- Cresswell, M., 1976. "The Semantics of Degree," in B. Partee (ed.), *Montague Grammar* (New York: Academic Press), 261–92.
- Davidson, D., 1967. "Truth and Meaning," *Synthese*, 17: 304–23. Repr. in *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press, 1984).
- 1978. "What Metaphors Mean," *Critical Inquiry*, 5: 31–47. Repr. in *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press, 1984), 245–64.
- 1986. "A Nice Derangement of Epitaphs," in E. Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (Oxford: Blackwell), 433–46.
- Deleuze, G., and F. Guattari, 1986. *Kafka: Toward a Minor Literature*, tr. D. Polan (Minneapolis: University of Minnesota Press).
- De Morgan, A., 1847. *Formal Logic or the Calculus of Inference, Necessary and Probable*. London: Tailor & Walton.
- DeRose, K., 1992. "Contextualism and Knowledge Attributions," *Philosophy and Phenomenology Research*, 52: 913–29.
- 1995. "Solving the Skeptical Problem," *Philosophical Review*, 104: 1–52.
- 1999. "Contextualism: Explanation and Defense," In John Greco and Ernest Sosa (eds), *The Blackwell Guide to Epistemology* (Malden, MA: Blackwell Publishers), 187–205.

## 168 BIBLIOGRAPHY

- Devitt, M., 1984. *Realism and Truth* (Oxford: Basil Blackwell).
- Eemeren, F. H. van (ed.), 2002. *Advances in Pragma-Dialectics* (Amsterdam: SicSat).
- and R. Grootendorst, 1992. *Argumentation, Communication, and Fallacies: A Pragma-Dialectical Perspective* (Hillsdale, NJ: Lawrence Erlbaum Associates).
- R. Grootendorst, F. S. Henkemans, J. A. Blair, R. H. Johnson, E. C. W. Krabbe, C. Plantin, D. N. Walton, C. A. Willard, J. Woods, and D. Zarefsky, 1996. *Fundamentals of Argumentation Theory: A Handbook of Historical Backgrounds and Contemporary Developments* (Mahwah, NJ: Lawrence Erlbaum Associates).
- Endicott, T., 2000. *Vagueness in the Law* (New York: Oxford University Press).
- Epstein, J., 2001. "Learning to be Thoughtless: Social Norms and Individual Computation," *Computational Economics*, 18: 9–24.
- Evans, G., 1978, "Can there be Vague Objects?" *Analysis*, 38: 208.
- Fara, D. G. [under pseud. of Delia Graff], 2000. "Shifting Sands: An Interest-Relative Theory of Vagueness," *Philosophical Topics*, 28: 45–81.
- 2008. "Profiling Interest Relativity," *Analysis*, 68: 326–35.
- 2012. "Literal Uses of Proper Names," manuscript, Princeton University. <<http://www.princeton.edu/~dfara/papers/fara-literal.pdf>>.
- Feldman, R. 2001. "Skeptical Problems, Contextualist Solutions," *Philosophical Studies*, 103: 61–85.
- Forbes, G., 1996. "Substitutivity and the Coherence of Quantifying in," *Philosophical Review*, 105: 337–72.
- Frege, G., 1956. "The Thought," tr. A. M. and Marcelle Quinton, *Mind*, 65: 289–311.
- Furnas, G., T. Landauer, L. Gomez, and S. Dumais, 1987. "The Vocabulary Problem in Human–System Communications," *Communications of the ACM*, 30: 964–71.
- Garrod, S., and A. Anderson, 1987. "Saying What you Mean in Dialogue: A Study in Conceptual and Semantic Co-ordination," *Cognition*, 27: 181–218.
- Gauker, C., 2002. *Words without Meanings* (Cambridge, MA: MIT Press).
- 2013. "Open Texture and Schematicity as Arguments for Non-Referential Semantics." Manuscript, Department of Philosophy, University of Cincinnati.
- Gentner, D. and Rattermann, M. J., 1991. "Language and the Career of Similarity," in S. A. Gelman, and J. P. Brynes (eds), *Perspectives on Language and Thought: Interrelations in Development* (London: Cambridge University Press), 225–77.

- Glucksberg, S. 2001. *Understanding Figurative Language: From Metaphors to Idioms* (Oxford: Oxford University Press).
- 2003. “The Psycholinguistics of Metaphor,” *Trends in Cognitive Science*, 7: 92–6.
- M. McGlone, and D. Manfredi, 1997. “Property Attribution in Metaphor Comprehension,” *Journal of Memory and Language*, 36: 50–67.
- Gray, A., 2012. *Names and Name-Bearing: An Essay on the Predicate View of Names*. Ph.D. dissertation, University of Chicago.
- Grice, P., 1989. “Logic and Conversation,” in Paul Grice, *Studies in the Way of Words* (Cambridge, MA: Harvard University Press), 22–40.
- Grim, P., P. St Denis, and T. Kokalis, 2003. “Learning to Communicate: The Emergence of Signaling in Spacialized Arrays of Neural Nets,” *Adaptive Behavior*, 10: 45–70.
- Grimshaw, J., 1990. *Argument Structure* (Cambridge, MA: MIT Press).
- Hale, K., and J. Keyser, 1987. “A View from the Middle,” *Lexicon Project Working Papers 10* (Cambridge, MA: Center for Cognitive Science, MIT).
- and — 1993. “Argument Structure,” in K. Hale and J. Keyser (eds), *The View from Building 20: Essays in Linguistics in Honor of Sylvain Bromberger* (Cambridge, MA: MIT Press), 53–109.
- Hauser, M. D., N. Chomsky, and W. T. Fitch, 2002. “The Faculty of Language: What is it, Who has it, and How did it Evolve?” *Science*, 298: 1569–79.
- Hawthorne, J., 2004. *Knowledge and Lotteries* (Oxford: Oxford University Press).
- Heck, R., 2002. “Do Demonstratives have Senses?,” *Philosophers' Imprint*, 2(2): <[www.philosophersimprint.org/002002](http://www.philosophersimprint.org/002002)>.
- Hesse, M., 1966. *Models and Analogies in Science* (South Bend, IN: University of Notre Dame Press).
- Higginbotham, J., 1985. “On Semantics,” *Linguistic Inquiry*, 16: 547–93.
- 1989a. “Knowledge of Reference,” in Alexander George (ed.), *Reflections on Chomsky* (Oxford: Blackwell), 153–74.
- 1989b. “Elucidations of Meaning,” *Linguistics and Philosophy*, 12: 465–18.
- Hills, D. 1997. “Aptness and Truth in Verbal Metaphor,” *Philosophical Topics*, 25: 117–53.
- Horn, L., 1989. *A Natural History of Negation* (Chicago: University of Chicago Press).
- and Ward, G. (eds), 2004. *The Handbook of Pragmatics* (Oxford: Blackwell).
- Kaplan, D., 1977. “Demonstratives,” manuscript UCLA. Repr. in J. Almog et al. (eds), *Themes from Kaplan* (Ithaca, NY: Cornell University Press, 1989), 481–566.

## 170 BIBLIOGRAPHY

- 1979. “On the Logic of Demonstratives,” *Journal of Philosophical Logic*, 8: 81–98.
- 1990. “Thoughts on Demonstratives,” in P. Yourgrau (ed.), *Demonstratives* (Oxford: Oxford University Press), 34–49.
- King, J., and J. Stanley, 2005. “Semantics, Pragmatics, and the Role of Semantic Content,” in Z. Szabo (ed.), *Semantics versus Pragmatics* (Oxford: Oxford University Press), 111–64.
- Kitcher, P., 1981. “Mathematical Rigor: Who Needs it?” *Noûs*, 15: 469–94.
- Kittay, E., 1990. *Metaphor: Its Cognitive Force and Linguistic Structure* (Oxford: Oxford University Press).
- Korta, K., and J. Perry, 2011. “Pragmatics,” in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2011 edition), URL = <<http://plato.stanford.edu/archives/sum2011/entries/pragmatics/>>.
- Kripke, S., 1979. “A Puzzle about Belief,” in A. Margalit (ed.), *Meaning and Use* (Dordrecht: D. Reidel), 239–83.
- 1980. *Naming and Necessity* (Cambridge, MA: Harvard University Press).
- 1982. *Wittgenstein on Rules and Private Language* (Cambridge, MA: Harvard University Press).
- Kurland, P., and G. Casper, 1975. *Landmark Briefs and Arguments of the Supreme Court of the United States: Constitutional Law* (Arlington, VA: University Publications of America), lxxv.
- Lakoff, G., 1973. “Hedges: A Study in Meaning Criteria and the Logic of Fuzzy Concepts,” *Journal of Philosophical Logic*, 2: 458–508.
- Lambert of Auxere, 1971. *Logica* (Summa Lamberti), ed. Franco Alesio (Florence: La Nuova Italia Editrice).
- Larson, R. K., 1998. “Events and Modification in Nominals,” in D. Strolovitch and A. Lawson (eds), *Proceedings from Semantics and Linguistic Theory (SALT) VIII* (Ithaca, NY: Cornell University Press.), 145–68.
- and P. Ludlow, 1993. “Interpreted Logical Forms,” *Synthese*, 96: 305–55.
- and G. Segal, 1995. *Knowledge of Meaning* (Cambridge, MA: MIT Press).
- Lascarides, A., and A. Copestake, 1998. “Pragmatics and Word Meaning,” *Journal of Linguistics*, 34: 55–105.
- Laserson, P., 1999. “Pragmatic Halos,” *Language*, 75: 522–52.
- 2005. “Context Dependence, Disagreement, and Predicates of Personal Taste,” *Linguistics and Philosophy*, 28: 643–86.
- Lepore, E., 1983. “What Model Theoretic Semantics Cannot Do?” *Synthese*, 54: 167–87.
- Lewis, D., 1969. *Convention* (Cambridge, MA: Harvard University Press).
- 1970. “General Semantics,” *Synthese*, 22: 18–67.

- 1979. “Score Keeping in a Language Game,” *Journal of Philosophical Logic*, 8: 339–59.
- 1982. “Logic for Equivocators,” *Noûs*, 16: 431–41.
- 1988. “Vague Identity: Evans Misunderstood,” *Analysis*, 48: 128–30.
- 1993. “Many, But Almost One,” in John Bacon, Keith Campbell, and Lloyd Reinhardt (eds), *Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong* (Cambridge: Cambridge University Press), 23–42.
- 1996. “Elusive Knowledge,” *Australasian Journal of Philosophy*, 74: 549–67.
- Ludlow, P., 1995. “The Logical Form of Determiners,” *Journal of Philosophical Logic*, 24: 47–69.
- 1999. “First Person Authority and Memory,” in M. de Caro (ed.), *Interpretations and Causes: New Perspectives on Donald Davidson’s Philosophy* (Dordrecht: Kluwer), 159–70.
- 2000. “Interpreted Logical Forms, Belief Attribution, and the Dynamic Lexicon,” in K. M. Jaszczolt (ed.), *Pragmatics of Propositional Attitude Reports* (Amsterdam: Elsevier Science), 31–42.
- 2002. “LF and Natural Logic,” in G. Preyer and G. Peter (eds), *Logical Form and Language* (Oxford: Oxford University Press), 132–68.
- 2005a. “A Note on Alleged Cases of Non-Sentential Speech,” in R. Eluguardo and R. Stainton (eds), *Elipsis and Non-Sentential Speech* (Dordrecht: Kluwer), 95–108.
- 2005b. “Contextualism and the New Linguistic Turn in Epistemology,” in G. Preyer and G. Peters (eds), *Contextualism in Philosophy* (Oxford: Oxford University Press), 95–108.
- 2006a. “Non-Sententialism and the FLN,” in T. Kawamura, Y. Suh, and R. K. Larson (eds), *Stony Brook University Occasional Papers in Linguistics, i. Intensionality and Sentential Complementation* (Stony Brook, NY: Dept. of Linguistics, Stony Brook University): <<http://www.linguistics.stony-brook.edu/sbopl>>.
- 2006b. “The Myth of Language,” *Croatian Journal of Philosophy*, 18: 385–400.
- 2008. “Cheap Contextualism,” *Noûs: Philosophical Issues*, 16: *Annual Supplement*, 104–29.
- 2011. *The Philosophy of Generative Linguistics*. Oxford: Oxford University Press.
- and S. Friend, 2003. “Disagreement and Deference: Is Diversity of Opinion a Precondition for Thought?” *Philosophical Perspectives*, 17: *Language and Philosophical Linguistics*, 115–39.
- MacFarlane, J., 2003. “Future Contingents and Relative Truth,” *Philosophical Quarterly*, 53: 321–36.

## 172 BIBLIOGRAPHY

- Mason, R., forthcoming. "Semantic Issues in the Social World," Ph.D. thesis, in preparation, Northwestern University.
- Metzing, C., and S. Brennan, 2003. "When Conceptual Pacts are Broken: Partner-Specific Effects on the Comprehension of Referring Expressions," *Journal of Memory and Language*, 49: 201–13.
- Mills, E., 2008. "The Egg and I: Conception, Identity, and Abortion," *Philosophical Review*, 117: 323–48.
- Neale, S., 1992. "Paul Grice and the Philosophy of Language," *Linguistics and Philosophy*, 15: 509–59.
- Neta, R., 2008. "How Cheap can you Get?" *Nous: Philosophical Issues*, 16: *Annual Supplement*, 130–42.
- Nirenberg, S., and V. Raskin, 1987. "The Subworld Concept Lexicon and the Lexicon Management System," *Computational Linguistics*, 13: 276–89.
- Ockham, W., 1951. *Summa Logicae, Pars Prima*, ed. P. Boehner (New York: St Bonaventura).
- Official Transcript: Proceedings Before the Supreme Court of the United States*, 1992. Washington, DC: Alderson Reporting.
- Perry, J., 1977. "Frege on Demonstratives," *Philosophical Review*, 86: 474–97.
- 2001. *Reference and Reflexivity* (Stanford, CA: CSLI Publications).
- Peter of Spain, 1972. *Tractatus Called Afterwards Summule Logicales*, ed. L. M. de Rijk (Assen: van Gorcum).
- Pinker, S., 1994. *The Language Instinct* (New York: William Morrow & Co.).
- Plunkett, D., and T. Sundell, 2012. "Disagreement and the Semantics of Normative and Evaluative Terms," manuscript, Dartmouth University and the University of Kentucky.
- Purver, M., 2004. "The Theory and Use of Clarification Requests in Dialogue," Ph.D. dissertation, University College London.
- Pustejovsky, J., 1995. *The Generative Lexicon* (Cambridge, MA: MIT Press).
- and S. Bergler (eds), 1991. *Lexical Semantics and Knowledge Representation* (Berlin: Springer-Verlag).
- Putnam, H., 1975. "The Meaning of 'Meaning,'" in K. Gunderson (ed.), *Language, Mind and Knowledge*, vii. *Minnesota Studies in the Philosophy of Science* (Minneapolis: University of Minnesota Press), 131–93.
- Quine, W. V. O., 1953. "Two Dogmas of Empiricism," in *From a Logical Point of View* (Cambridge, MA: Harvard University Press), 20–46.
- 1980. "Reference and Modality," in *From a Logical Point of View* (Cambridge, MA: Harvard University Press), 139–59.
- Raffman, D., 1996. "Vagueness and Context-Sensitivity," *Philosophical Studies*, 81: 175–92.

- “Rape.” 1992. *American Jurisprudence* (Rochester, NY: Lawyers Cooperative), 2nd edn, lxxv [cumulative supplement]. 88–149.
- “Rape and Battery: Between Husband and Wife,” 1954. *Stanford Law Review*, 6: 719–28.
- Rayo, A., 2013. “A Plea for Semantic Localism,” *Noûs*, 47: 647–79.
- Recanati, F., 1993. *Direct Reference* (Oxford: Blackwell Press).
- 2004. *Literal Meaning* (Cambridge: Cambridge University Press).
- 2007. *Perspectival Thought* (Oxford: Oxford University Press).
- 2010. *Truth-Conditional Pragmatics* (Oxford: Oxford University Press).
- Richard, M., 1990. *Propositional Attitudes* (Cambridge: Cambridge University Press).
- Russell, D., 1982. *Rape in Marriage*. New York: Macmillan.
- Sánchez Valencia, V., 1994. “Monotonicity in Medieval Logic,” in A. de Boer and H. de Swart (eds), *Language and Cognition 4* (Groningen: Yearbook 1994 of the Research Group for Theoretical and Experimental Linguistics at the University of Groningen).
- Scalia, A., 1997. *A Matter of Interpretation* (Princeton: Princeton University Press).
- Schegloff, E., 2007. “Categories in Action: Person Reference and Membership Categorization,” *Discourse Studies*, 9: 433–61.
- Schiappa, E., 2003. *Defining Reality: Definitions and the Politics of Meaning* (Carbondale, IL: Southern Illinois University Press).
- Schober, M., and H. Clark, 1989. “Understanding by Addressees and Overhearers,” *Cognitive Psychology*, 21: 211–32. Repr. in H. Clark (ed.), *Arenas of Language Use* (Chicago: CSLI Publications and University of Chicago Press, 1992), ch. 6.
- Shapiro, S., 2006. *Vagueness in Context* (Oxford: Oxford University Press).
- Sidnell, J., 2010. *Conversation Analysis: An Introduction* (New York: Wiley-Blackwell).
- Silberstang, E., 1972. “Rape,” *American Jurisprudence*, 2nd edn (Rochester, NY: Lawyers Cooperative), lxxv. 757–838.
- Skyrms, B., 2004. *The Stag Hunt and the Evolution of Social Structure* (Cambridge: Cambridge University Press).
- 2010. *Signals: Evolution, Learning and Information* (Oxford: Oxford University Press).
- Soames, S., 1999. *Understanding Truth* (Oxford: Oxford University Press).
- Sommers, F., 1970. “The Calculus of Terms,” *Mind*, 79: 1–39.
- Sperber, D., and D. Wilson, 1986. *Relevance: Communication and Cognition* (Oxford: Blackwell).

## 174 BIBLIOGRAPHY

- Stainton, R., 2005. "In Defense of Non-Sentential Assertion," in Z. Szabo (ed.), *Semantics versus Pragmatics* (Oxford: Oxford University Press), 383–457.
- Stanley, J., 2000. "Context and Logical Form," *Linguistics and Philosophy*, 23: 391–434.
- 2002a, "Making it Articulated," *Mind and Language*, 17: 149–68.
- 2002b. "Nominal Restriction," in G. Peters and G. Preyer (eds), *Logical Form and Language* (Oxford: Oxford University Press), 365–88.
- 2003. "Context, Interest Relativity and the Sorites," *Analysis*, 63: 269–80.
- 2005. *Knowledge and Practical Interests* (Oxford: Oxford University Press).
- and Z. Szabo, 2000. "On Quantifier Domain Restriction," *Mind and Language*, 15: 219–61
- Stern, J., 2000. *Metaphor in Context* (Cambridge, MA: MIT Press).
- Stewart, I., 1998. *Life's Other Secret* (New York: John Wiley & Sons).
- Stine, G., 1999. "Skepticism, Relevant Alternatives, and Deductive Closure," in K. DeRose and T. Warfield (eds), *Skepticism: A Contemporary Reader* (Oxford: Oxford University Press), 145–55.
- Strogatz, S., 2003. *Sync: The Emerging Science of Spontaneous Order* (New York: Hyperion).
- Tappenden, J., 1999. "Negation, Denial and Language Change in Philosophical Logic," in D. Gabbay and H. Wansing (eds), *What is Negation?* (Dordrecht: Kluwer), 261–98.
- Taylor, K., 1995. "Meaning, Reference, and Cognitive Significance," *Mind and Language*, 10: 129–80.
- Teller, P., 2010. "Modeling Truth," manuscript, University of California Davis.
- "To Have and to Hold: The Marital Rape Exemption and the Fourteenth Amendment," 1986. *Harvard Law Review*, 99: 1255–73.
- Travis, C., 1985. "On What is Strictly Speaking True," *Canadian Journal of Philosophy*, 15: 187–229.
- 1989. *The Uses of Sense* (Oxford: Oxford University Press).
- 1996. "Meaning's Role in Truth," *Mind*, 100: 451–66.
- Tribe, L., 1992. *Abortion: The Clash of Absolutes* (New York: Norton).
- Tyson, N., 2009. *The Pluto Files* (New York: Norton & Co.).
- Unger, P., 1979. "I do Not Exist," in G. F. Macdonald (ed.), *Perception and Identity: Essays Presented to A. J. Ayer with his Replies to Them* (New York: Macmillan), 235–51.
- 1984. *Philosophical Relativity* (Minneapolis: University of Minnesota Press).
- 1986. "The Cone Model of Knowledge," *Philosophical Topics*, 14: 125–78.

- van Deemter, K. and S. Peters (eds), 1997. *Ambiguity and Underspecification* (Stanford, CA: CSLI Publications).
- Van Inwagen, P., 1990. *Material Beings* (Ithaca, NY: Cornell University Press).
- Venuti, L., 1995. *The Translator's Invisibility: A History of Translation* (London: Routledge).
- Watts, D., 1998. *Small Worlds: The Dynamics of Networks between Order and Randomness* (Princeton: Princeton University Press).
- Weintraub, D., 2007. *Is Pluto a Planet?* (Princeton: Princeton University Press).
- Wilkes-Gibbs, D., and H. Clark, 1992. "Coordinating Beliefs in Conversation," *Journal of Memory and Cognition*, 31: 183–94.
- Wilson, D., and R. Carston, 2007. "A Unitary Approach to Lexical Pragmatics: Relevance, Inference and Ad Hoc Concepts," in N. Burton-Roberts (ed.), *Pragmatics* (London: Palgrave), 230–59.
- Wittgenstein, L., 1956. *Remarks on the Foundations of Mathematics*, tr. G. E. M. Anscombe (Cambridge, MA: MIT Press).
- 1958. *Philosophical Investigations*, tr. G. E. M. Anscombe (New York: Macmillan).
- 1969. *On Certainty* (New York: Harper & Row).
- Woodger, J., 1937. *The Axiomatic Method in Biology* (Cambridge: Cambridge University Press).
- 1939. *The Technique of Theory Construction: International Encyclopedia of Unified Science* (Chicago: University of Chicago Press), ii/5.
- Zadeh, L., 1965. "Fuzzy Sets," *Information and Control*, 8: 338–53.
- Zivanovic, S. (2002). *Kvanti kacija v naravnem cloveskem jeziku*. Diplomsko delo, Univerza v Ljubljani, Fakulteta za matematiko in ziko. An updated English version is also available from [saso.zivanovic@guest.arnes.si](mailto:saso.zivanovic@guest.arnes.si).

