

# 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 to 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 imbroglions 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 ‘hacktivist’ 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.alexao'Brien.com/secondsight/letter\\_to\\_david\\_carr\\_ravi\\_somaiya\\_and\\_the\\_new\\_york\\_times.html](http://www.alexao'Brien.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 (o.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.

## 46 NORMS OF WORD MEANING LITIGATION

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.

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 supposed 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.

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.

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 Janes’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.