


OXFORD

# Living Words

Meaning Underdetermination  
and the Dynamic Lexicon



PETER LUDLOW

## 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: 2014932276

ISBN 978-0-19-871205-3

Printed and bound by  
CPI Group (UK) Ltd, Croydon, CRO 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 litigates 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 Recanatì, 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 Beihai 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,

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. A portion of section 1.2 draws on my article “What Is a Hacktivist?” *New York Times Opinionator Blog*, January 13, 2013. 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>Keyword Index</i>	177
<i>Names Index</i>	187
<i>Organizations Index</i>	191



# 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 fifty 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.

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.



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

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

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

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

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

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

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

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

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

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),

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 evaluate 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).