


JODY AZZOUNI

SEMANTIC PERCEPTION

*How the Illusion of a Common Language
Arises and Persists*



SEMANTIC PERCEPTION

This page intentionally left blank

Semantic Perception

*How the Illusion of a Common Language
Arises and Persists*

Jody Azzouni

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

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 New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trademark of Oxford University Press in the UK and certain other countries.

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

© Oxford University Press 2013

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 license, or under terms agreed with the appropriate reproduction rights organization. Inquiries 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.

Library of Congress Cataloging-in-Publication Data

Azzouni, Jody.

Semantic perception: how the illusion of a common language arises and persists / Jody Azzouni.

p. cm

Includes bibliographical references (p.) and index.

ISBN 978-0-19-996740-7 (hardback: alk. paper) — ISBN 978-0-19-996741-4 (updf)

1. Semantics (Philosophy) 2. Semantics—Methodology. 3. Meaning (Philosophy)

4. Perception (Philosophy) 5. Nonverbal communication—Philosophy. I. Title.

B840.A99 2013

149'.94—dc23

2012036334

1 3 5 7 9 8 6 4 2

Printed in the United States of America
on acid-free paper

Contents

Acknowledgments vii

General Introduction 1

0 Methodological Preliminaries 14

1 The Experience of Understanding Expressions 47

2 Illusions of Meaning 82

3 The Phenomenological *What Is Said* 113

4 The Experienced Distinction between *What Is Said* and
Implicated Content 143

First Methodological Interlude: Special Sciences and
Evidence 167

5 Strict Content 179

6 Truth Content 206

Second Methodological Interlude 245

7 The Use and Misuse of Communicative Intentions:
Grand-Style Neo-Griceanism 265

8 Opportunistic Applications of Posited Communicative
Intentions 290

9 Artificial Extensions of Natural Language 321

General Conclusion 348

Bibliography 355

Index 365

This page intentionally left blank

Acknowledgments

I've been working on this book for quite a few years: I started it back in 2004 but dropped it to write a different book in 2005 and took it up again only to drop it a second time to write still another book in 2008. (My book projects don't usually survive being dropped repeatedly this way: I've got an awful lot of unfinished books.)

I've accrued debts along the way to finishing this one. My thanks, first of all, to George Smith, for supplying what was virtually a line-by-line commentary on an earlier version of the book. His suggestions—about content, about organization, even about syntax—were invaluable. Others were helpful too. My thanks to both John Collins and Brian Epstein, for numerous substantial and valuable suggestions, and to Jeff McConnell and Gary Ostertag, for a number of useful conversations on the various topics and contents of the book. Finally, I taught material from the book twice at Tufts to our excellent students; my thanks to them as well.

This page intentionally left blank

SEMANTIC PERCEPTION

This page intentionally left blank

General Introduction

I want to state as simply and as plainly as I can the major thesis of this book. It's this: human beings *involuntarily* experience certain physical items, certain products of human action, and certain human actions themselves, as having monadic meaning-properties: for example, as possessing meanings, as referring, or as having (or being capable of having) truth values—thus the main title of the book, *Semantic Perception*. My thesis is that we (human beings) involuntarily see uttered words, among other things, as possessing certain monadic meaning-properties, and that we involuntarily see uttered sentences as possessing other (but related) monadic meaning-properties.

By “monadic meaning-properties” I mean that we experience these properties as properties *of* uttered words and sentences similar to how we perceive ordinary objects to have as monadic properties shape and color (but not location). We don't experience the meaning-properties of uttered words and sentences (for example) as due to background conventions or regularities or as relations to context, nor do we experience them the way that we experience codes. We *don't* see them as relations (between uttered words and sentences and us) that are due to the interactional effects of the items in question and our minds, or aspects of our minds; we don't experience these meaning-properties as derivative from the expectations and intentions of the producers of these items (or as due to the mutual actions of speakers and the audience).

This doesn't mean that we don't, during conversation or during reading events, recognize the expectations or intentions of speakers or writers (or those of our own). We always recognize, for example, that the sentences

uttered in conversations are uttered *intentionally*. It does mean, however, that we *experience* the meaning-properties—the *meanings*—of the produced items as *independent* of the speaker’s intentions in exactly the same way that we experience an object’s shape as independent of its color.

I call the view that human beings see a large class of physical objects and human actions (such as pointing) as possessing monadic meaning-properties the *semantic perception view*.¹

I don’t claim that when we see meaningful physical items, such as the words on this page, or when we hear people speak that we see these physical instantiations of words and sentences to have meaning-properties the way that we *see* a book to have a red cover. That’s why I’ll often fall back on the broader word “experience” rather than “see” or “perceive.” Nevertheless, “see,” “perceive,” and their cognates are usually understood broadly enough all on their own that “semantic perception” shouldn’t mislead.

It has seemed to many philosophers that we have only very few viable choices for systematically explaining the apparent meaning-properties that public (and sometimes privately understood) physical events—such as the sounds we make when we are speaking to one another or inscriptions on paper or on electronic media—seem to us to have. One possibility is that such properties *really are* ones that such physical events have—derivatively perhaps, by virtue of being instances of certain *types*. Those inclined toward certain social constructivist views might want to claim that language entities *really have* the meaning-properties they seem to have because, of course, *we* (collectively) designed them—the socially constructed types—to have such properties. Similarly, a sculpture really has the property of being the sculpture of someone because the sculptor has endowed it with that property.

It’s widely claimed that certain designed objects are items with which we have endowed meaning-properties; it’s widely claimed that such designed objects have “derived intentionality”² (as it’s so often put in the literature). This, however, is metaphor that must be unpacked: the mechanisms by which we endow (and can endow) designed objects with intentionality need to be explained. More strongly, the nature of the mechanisms by which it’s possible for us to endow an object with intentionality needs to be explained.

Many philosophers think that only one family of approaches to language phenomena can meet this demand. This is to treat the meaning-properties of public-language events and objects as derived or *projected* properties that are based in some way on human psychology. The communication events between humans are thus treated as the central genuine phenomenon of language. Two or more people during a language event have intentions to communicate and to understand one another, and a characterization of how they use language events and objects to manage this must be derived from

1. The emerging cognitive science literature on mirror neurons suggests the thesis that we are *not* the only animal that experiences events in the world—specifically the actions of other animals—as *meaningful*. See the relevant articles in Hurley and Chater (2005).

2. See Dennett (1987), and the references therein.

the propositional-attitude psychology of the participants. That is, it must be derived from their background mutual knowledge and understanding (for example, of public conventions) as well as from their mutually perceived communicative intentions, expectations, and so on. I mean, of course, to allude to the huge family of approaches to public language that originated from Paul Grice's extremely influential work.

The project of reducing our understanding and practice of language (via mutual understanding and various communicative intentions) to propositional-attitude psychology has seemed to many philosophers to be the only philosophical game in town when it comes to explaining how public-language entities are invested with meaning-properties. According to this family of intention-based approaches, the derivative nature of the meaning-properties of actual public languages is taken for granted. There are only the individual psychologies of the participants, and everything apparently involved in the meaningfulness of public languages—including the apparently successful communication acts based on the comprehension of the utterances of those languages and the apparent presence of complex public conventions that describe meaning regularities—must be characterized on the narrow basis of individual psychology and behavior.

Many neo-Griceans³ want to claim that “public languages” don't contain any objects or events with monadic meaning-properties. Rather, what public languages are is this: they are classes of physical entities (or events)—used in common by a population—that are endowed with meanings by virtue of a set of conventions that have been mutually adopted, and are mutually known to be adopted, by that population. That is, successful communication events occur (by means of these classes of commonly-used entities and events) via mutually recognized communicative intentions the presence of which is indicated, at least in part, by the deliberate use of these classes of commonly-used entities.

So, on these views, the language tokens of various sorts that we use to communicate are themselves meaning inert. Their meaning-properties are derived entirely from the intentions and mutual knowledge of their users in much the way that two people playing chess understand each other to be employing chess-property-inert physical objects that are endowed only derivatively with specific “chess properties” by the mutual understanding and by the knowledge of background conventions of the two players. I'll discuss such views in chapter 7.

There is an overlooked alternative approach to public language, however, that agrees with Gricean and neo-Gricean approaches that the apparent meaning-properties of public language entities must be derived from human psychology. The disagreement, however, is over exactly how this derivation occurs. The disagreement is over *which* psychological traits of human beings

3. In this book I'll generally use “neo-Gricean” more broadly than some do. I include, for example, relevance theorists among neo-Griceans.

are relevant to understanding the effortless communication events we engage in. According to both views, none of the physical objects utilized by humans in their communication activities—written words, spoken sentences, and so on—*actually have* any meaning-properties, are *actually* anything more than pure physical objects and events. The semantic perception view differs from Gricean and neo-Gricean ones, however, in claiming that we involuntarily *experience* written words, spoken sentences, and so on, as having monadic meaning-properties. We *experience* them, for example, as items that refer, and that are meaningful.

As with Gricean and neo-Gricean intention-based approaches to public language, all public-language phenomena are still to be explained in terms of individual psychological phenomena. In sharp contrast to those approaches, no *systematic* deployment of communicative intentions and expectations is needed for this task, nor are notions of mutual knowledge (used by Griceans to ground public conventions) called for. This is because if two people involuntarily experience an uttered sentence as monadically meaning something, then that perceived meaning is (of course) the default experience of what that uttered sentence means. That is, the uttered sentence is experienced as meaning what it's perceived to mean by virtue of its own meaning-properties—unless the utterer is seen as deliberately *tampering with* that meaning (for example, by sarcasm).

My primary aim in this book isn't to refute Gricean and neo-Gricean approaches once and for all. Griceanism is too widespread and its numerous proponents are far too varied in their individual approaches.⁴ My primary aim is more modest: to make a case for the viability of the alternative semantic perception view. While doing so in the course of this book, I'll repeatedly respond to specific Gricean and neo-Gricean doctrines and along the way modify much of the semantic-pragmatic apparatus loved by Griceans, such as *what is said* or *implicatures*. My aim is to show how Gricean assumptions about the centrality of mutual knowledge and communicative intentions to the phenomena of perceived meaning-properties badly distort ordinary folk-psychological attributions of intentions, beliefs, and expectations as well as ordinary intuitions about what is said, what is implied but not said, and so on. *It's an important aspect of the evidence for the semantic perception view that it explains and sustains the ordinary phenomenology of the experience of understanding language, whereas Gricean and neo-Gricean views instead consistently distort or attempt to explain away this phenomenology.* This is the primary line of argument throughout this book.

I'm *not* claiming that the perception of monadic meaning-properties utterly supplants recognition of communicative expectations and intentions. As the truism that we utter statements *intentionally* makes clear, there is a large class of language events that we perceive as involving intentions and expectations. One way of putting the difference between the semantic perception view

4. And besides, who needs to make *that many* enemies all at once?

and its Gricean and neo-Gricean opponents is that on the latter views, intentions and expectations of speakers perceivably play a constitutive role in our experience of meaning. On the semantic perception view, such intentions and expectations are perceived to play an ancillary role, one in which the perceived monadic meaning-properties of uttered expressions are *presupposed*.

I should note two other theses, of a more broadly metaphysical nature that I hold, and that have a presence in this book. The first is nominalism. On my view, there are no properties, no universals, no mathematical objects, and no types. I've long argued (for example, in Azzouni 2004b, 2010b) that this nominalism leaves intact our ordinary ways of talking about properties, universals, mathematical objects, or types. Indeed, I've long argued that our true descriptions and theories of the world and of phenomena within the world *must* encompass quantification over what doesn't exist. The nominalist position I espouse, that is, is a purely metaphysical one that places no constraints on the forms our descriptions and theories of the world must take. I will say more about this presupposed nominalism in section 0.2.

Acceptance of the semantic perception view doesn't require an acceptance of nominalism in my sense. The views do fit nicely together, however, because the nominalist can cleanly describe the experiences of speaker-hearers by sentences that quantify over things that don't exist. This aspect of the semantic perception view, so I claim, isn't unique to it: quantifying over what doesn't exist shows up widely in the sciences as I discuss in Azzouni (2010a, chapter 4), and as I'll also discuss in section 0.2.

A consequence of this nominalism, however, is that there is no language, such as English, that we speak in common. Any such language would have to be composed of types, and there are no types. What languages are, ontologically speaking—*all they are*—are specific communication events: actions taken by people during which they produce noises and experience an understanding of one another. In addition, of course (as this book makes plain), there are numerous other artifacts, physical objects of various sorts, that we treat as comprehensible instances of language. And of course, the people engaging in language transactions are in various relevant psychological states—where the states in question needn't be ones they are conscious of.

The second thesis is this: there are no physical objects in the world with meaning-properties of any kind. No object, that is, has the property of meaning some other thing in the way that we experience words to so mean what they refer to. Nothing real, that is, refers; nothing real monadically *means* anything. No more does a mirror image or a picture of something refer or otherwise indicate what it seems to refer to. The relationship between a mirror and what it reflects is a pure causal relation; the same corresponding claim is true of a picture. The *reference relationship* between a word and what it refers to is projected upon these items by persons who so experience words as so referring to things; any meaning relationship between something and something else can only be one that's imposed upon these things by an experiencing person.

As with my first thesis, the denial of this claim is compatible with the semantic perception view. After all, it could be (it could have been?) that we experience public-language entities as having monadic meaning-properties *that they actually have*. I think this is evidently false (and I won't be engaging, in this book anyway, with those who think otherwise).

One immediate qualification of this second thesis is needed. I'll be leaving aside the rather important issue of human beings, their brains, and perhaps other intelligent creatures and their relevant organs. I'm leaving aside the issue of "original intentionality." Human beings, on my view, are also (and nothing more than) physical objects. Having said this, I hasten to add that I don't mean to legislate on the question of whether the human mind itself has, or has contents with, meaning-properties. The second thesis, therefore, is restricted in its scope only to human *products* of various sorts, where the extension of "product" is understood to include perceived meaningful actions.

Before turning to an overview of the contents of this book, I want to briefly mention two factors in recent intellectual history that I believe prevented the semantic perception view from coming to light earlier. One barrier for the semantic perception view is that it *does* attribute a widespread and psychologically involuntary misapprehension to ordinary people. Because of how they involuntarily experience language phenomena, they are impelled to think the words and sentences of their language have an interlocked system of properties that words and sentences don't have (that *nothing* has). That this is a misconception will be shown step by step throughout the course of this book. Indeed, the semantic perception view *just about* attributes a systematic *hallucination* to the average speaker-hearer—indeed to all of us.

"Hallucination," however, puts the matter far too strongly. The problem is with the word itself. We understand hallucinations, as Austin pointed out long ago, to be experiences of nothing at all. Our standard understanding of an hallucination is as an experience of an object where no object is (located). The semantic perception view doesn't imply that we're having experiences of this sort: the view is that we experience objects and events, that we otherwise are aware of, as endowed with monadic properties that they don't have. The ways we ordinarily speak of *illusions* doesn't quite fit the experiences I'm speaking of either (although I've capitulated in using "illusion" in the subtitle of this book: how the illusion of a common language arises and persists). This is because "illusions" aren't normally applied to a network of experiences as systematic and interconnected as our experience of language entities is. Illusions are understood to be illusions about specific objects, but relatively unsystematic ones. There is no word that combines the systematic features that hallucinations are understood to be capable of along with the fact that, as with illusions, it's actual objects that are so misperceived.

A solution, I suppose, would be the coining of new words to do the job needed—for example, "illusionate" and "illusionation." An *illusionation*, it could be said, is the involuntary systematic experience of an object or class of objects as having properties that they in fact don't have. We *illusionate*

physical objects of various sorts as having meaning-properties—for example, we illusionate them as monadically referring to certain things and not to other things.⁵

I considered writing the entire book utilizing this new jargon. I suspect almost everyone reading this will be pleased to learn that in the end I decided not to do so.

A second historical factor preventing the emergence of the semantic perception view was the widespread lesson drawn by most philosophers from the “Quine-Putnam” indispensability thesis. If we *must* speak of a kind of entity, existentially quantify over it, then we have no choice but to be ontologically committed to it. Claiming, therefore, that we involuntarily perceive types or tokens with monadic meaning-properties, and claiming that as a result we *must* characterize language and communication in terms of such entities, would have struck an entire earlier generation of philosophers as simply a concession that such things *exist*. If one denies the existence of physical entities with monadic meaning-properties, then only by rejecting Quine’s criterion for what a discourse is committed to is one enabled to entertain the semantic perception view as a viable alternative to Gricean and neo-Gricean approaches.

One last point before I turn to sketching out the contents of this book. There is an important layering to how we experience the properties of words and sentences. It isn’t merely that we involuntarily experience sentences and words as having meaning-properties; it’s also that the meaning-properties that we experience sentences and words to have are affected by factors that we are unconscious of the impact of. I develop one example of this phenomenon at length: our experience of what is said by an utterance is affected by various contextual factors. For the most part we are so unaware of the role of these contextual factors that we experience the meaning-properties of sentences to be ones that are independent of *anything*. So it’s not merely that our minds—as it were—project meaning-properties onto otherwise meaning-inert entities and events; it’s that the meaning-properties we’re conscious of aren’t the same as the meaning-properties that our minds (our language faculties) utilize to enable us to understand and express meaningful statements. There’s a “disconnect” between our consciously-accessible experience of the meaning-properties of language objects and events and the meaning-properties that our subpersonal language faculties project onto those objects and events. Both are “projections” because statement events and statement objects in themselves have no meaning-properties whatsoever, but the projections involved are different ones, projected by apparently quite different “faculties” of mind. This rather striking and disconcerting phenomenon is illustrated several times and in several chapters.

5. What strengthens my temptation to introduce new jargon in this way is my view that we are constantly having illusionations: we constantly perceive objects to have properties that in fact they don’t have. That is, illusionations aren’t restricted to language entities. But that’s a topic for another time and another book.

I'll finish this general introduction by sketching the contents of the book itself. Chapter 0 provides important methodological preliminaries. In 0.1, I describe the way that the nonphilosopher naturally confuses types and tokens. This is by way of explaining some of the psychological sources for Platonism. I next turn to indispensability issues in 0.2. The traditional literature on the (Quine-Putnam) indispensability thesis understands this phenomenon in a purely language-based way—as a matter of how our scientific and commonsense theories must quantify over and refer to Platonic entities. I suggest an additional psychological version of indispensability, one that manifests itself in how we must *experience* certain things, for example during language transactions. Section 0.3 takes up some of the debate in the literature about the existence of public languages. The purpose is to situate the semantic perception view in relation to these earlier discussions. Last, in 0.4, I take up the issue of “phenomenology.” My purpose is to explain the methodological value of correctly characterizing what's involved in our experience of understanding language and perceiving meanings. One important methodological value, that I stress here and later in the book, is that “intuitions” are a central tool in this area of philosophy of language, and in philosophy of language generally. It's only by systematically studying the phenomenology of our language experience that we can begin to evaluate the appropriateness of such intuitions as data.

The overall aim of the first six chapters is to present the phenomenology of language perception in some detail. It's to show that out of the confused (and oftentimes bizarre) experience of the understanding of the expressions of natural languages, the notion of a common public language inexorably arises. This is a notion accompanied by the idea that we can share truths, and pool them together for common use. I claim and try to show that it's on the basis of this notion of a common public language that we have constructed our notions—and actual institutions—of science and mathematics.

To this end, in chapter 1, I first attempt to characterize the intricate ways that types and tokens apparently weave through our experience of the understanding of uttered expressions. When we attribute truth or falsity to something, sentence-type descriptions seem to be needed in the background of our attribution or the attribution will be heard as incoherent. Furthermore, these types seem to so dominate our conception of meaningful tokens that we experience all such tokens as “similar” to one another. On the other hand, our perception of such types is individuated in certain respects by the tokens that instantiate them: we perceive such types to have properties due to their being physically localized to where such tokens are. This has a profound effect on the nature of how we experience the metaphysics, as it were, of “what is said” when someone utters an expression.

I turn in section 1.3 to a description of how the apparent interanimation of types and tokens affects our experience of the references of our words. The main lesson of the section is how oblivious we are to the actual pattern of referential usage that our words instantiate. Here is one way that our

own obliviousness (itself) is obscure to us—how, that is, we're ignorant of our ignorance. We apparently experience the references of the tokens of our terms to be simply due to the types they instantiate, yet we have no conscious access to the mechanisms (if any) by which such types govern their tokens. One interesting side-effect of this is a systematic mismatch between nonphilosophers' "ontological views"—their descriptions of what their terms refer to—and the actual referential pattern of usage of those terms.

Section 1.4 turns to a concern with how speaker-hearers' impressions of utterances/sentences constrain their views of the mechanisms by which such items are meaningful to them. They have a "what you see is what you get" picture of sentences: their properties and their constituents are exactly what speaker-hearers see them to be. As a result, speaker-hearers are impelled to see "what is said" by a straightforwardly uttered expression to be the speaker's thought so conveyed without having any corresponding impression of how this happens except vis-à-vis the visible elements in those expressions. A theme first arises here that recurs throughout the book: contextual elements aren't seen as playing a constitutive role in what is said by an expression uttered. Indeed, gestures toward contextual elements (*this* room) or toward psychological factors (intentions to refer to *this* room) are seen as going beyond the expression uttered and, similarly, beyond what's been said by that uttered expression.

A second property that speaker-hearers attribute to uttered expressions is that the meanings of such are monadic properties of them. Furthermore (i) this experience is involuntary and automatic, and (ii) nothing in this experience of understanding uttered expressions enables speaker-hearers to recognize the constitutive or causal role of contextual factors or intentions that are contributing to the meanings of uttered expressions.

In Chapter 2, I continue my investigation of the experience of understanding uttered and inscribed expressions. First, I hypothesize that the source of nonprofessionals' systematic fumbblings over types and tokens in language is the involuntary (but largely subliminal) injection of a quite different notion—that of a functionally and physically propertied *object*. The suggestion is that speaker-hearers automatically experience and talk about sentences and words the way they experience and talk about tools, like screwdrivers and hammers.

In the sections that follow, I present a number of "meaning illusions." These are cases where the experience of perceiving shapes as *meaningful* is irresistible—even given the background knowledge that the shapes one experiences as meaningful are the results of nonintentional accidents (such as erosion). Examples include cases of involuntary saturation, where pronouns, say, are experienced automatically as referring to certain accidentally salient individuals, or where ambiguous terms are disambiguated by contextual factors that just happen to be present. The accidental shapes in question aren't only seen as meaningful. In the right circumstances, they are also seen as truth apt. Meaning illusions include cases where we experience pronouns or names as

successfully referring to someone or other, although (and this is part of the experience) we don't know who.

An interesting aspect of how we experience meaning is also discussed further in this chapter: How alternative readings of an expression (that are possible because of different contextual factors or because of ambiguity) are invisible in the moment of our experiencing the particular meaning we experience an utterance of that expression to have. This “phenomenologically compartmentalized” aspect of the experience of understanding uttered expressions is due, in large measure, to the invisibility of the roles of context and speaker-meaning in *what is said*.

Strikingly, our involuntary experience of meaning *as* a property of the uttered expression is coupled with a *simultaneous* recognition (experience) of the pure physicality of that uttered expression (for example, as sheer ink marks on paper): the recognition, that is, of the *projective nature* of our experience of that meaning.

The chapter ends with a discussion of why tempting inferentialist construals of the above phenomena should be resisted.

Chapters 3 and 4 take up the distinction between what is said and what is implicated but not said. The aim is to distinguish purely phenomenological versions of these notions (versions with properties to be characterized entirely in terms of the experience of speaker-hearers) from the various analogous theoretically infiltrated notions currently rampant in the Gricean and neo-Gricean literature. Of course, no “pure phenomenological” notion either of what is said or of what is implicated but not said will be free of implicit generalizations or of implicit inductions of one kind or another on the part of the nonprofessional speaker-hearer. But these will result from aspects of their experiences—from what *they* are conscious of. The notions *will* be free of opportunistic theoretical assumptions of language theorists that massage the evidence (of what the speaker-hearer is conscious of) so that it better conforms to one or another antecedent theory of language, understanding, or communication.

An important issue is raised in section 4.1 that I'd like to draw your attention to now; just as speaker-hearers aren't conscious of “speech acts,” they are similarly unconscious of the presence of context (mutual, or other) in conversational transactions. At best (for example, in the recognition of phenomenologically available implicatures) speaker-hearers are aware of what they *take* to be the mutual awareness of salient items in their immediate neighborhood and similarly what they *take* to be mutually held items of knowledge.

In chapter 5, I turn to the notion of strict or literal content. In contrast to the notions of what is said and what is implicated but not said that are automatic experiences had by speaker-hearers when they are in the presence of utterances or written inscriptions, the ordinary notion of what an expression says strictly or literally is one that speaker-hearers rely on when they take a “second look” at an expression. In interesting ways, not shared by what is

said and what is implicated, what is strictly said is theoretically infiltrated or can easily be made to be so infiltrated.

Both the nonprofessional speaker-hearer and the early theorist of language who relied directly on the intuitions of that speaker-hearer regard what is said as almost always nearly identical to what is literally or strictly said by the expression used. This is because of the aforementioned invisibility of contextual and intentional effects on what is said, as well as the general invisibility of the alternative interpretations our expressions are open to. To put the point in a general way: it's because of the involuntary phenomenological compartmentalization of our experience of understanding expressions. That is, it's intuitive to regard the literal meaning of an expression to be usually what's communicated in the speech situation. This, in turn, invites the impression that expressions are self-standing vehicles of meaning with stable properties that we can exchange with one another in communication situations, and that we can carry away from one communication situation and introduce in later quite different communication situations without having to change them in any way.

In chapter 6, the benefits are reaped of the various aspects of the phenomenology of the understanding of uttered and inscribed expressions described in earlier chapters. Our truth sharing practices as we see them are characterized and contrasted with what in certain respects they're really like. It's explained how we can see ourselves passing around truths via truth vehicles that don't change their properties as they change hands—how the expressions we utter can be seen as remaining the same as we take ourselves to repeat them in new contexts. One important tool that enables this impression simultaneously gives rise to Donnellan's referential-attributive distinction, and this explains why the latter distinction is so intuitively accessible.

In section 6.4, I take up a description of the modest grasp the nonprofessional speaker-hearer has of compositionality. It's a grasp that collapses under even moderate inspection. It's sufficient, however, to give speaker-hearers the impression that their sentences have the meanings they have because of their visible subcomponents and that exceptions to this—idioms—are relatively rare. Nonprofessional speaker-hearers, of course, have no idea how little they grasp about the rules by which the meanings of sentences are derived from the meanings of their parts.

Discussion of the parts of sentences—for example, words—is naturally followed by (in section 6.5) the speaker-hearer's impression of the concepts that accompany those words. Included is a discussion of the reasons for the nonprofessional speaker-hearer's unawareness of the limitations in the applications of words—why, that is, bivalence is so intuitive despite its nearly complete absence from natural languages.

These various elements—arising in different ways from the psychological experience of the speaker-hearer (when in communication situations)—conspire together to create a very neat, but utterly false, picture of the natural languages that we take ourselves to speak. It's one that has taken us over a century of sustained professional work to escape from.

In chapters 7 and 8, I turn to a more sustained attack on Gricean and neo-Gricean intention-based approaches to semantics. These approaches are notoriously riddled with counterexamples and problems. My aim is to rehearse some of these and illuminate them by showing how they are due to distorted descriptions of the ordinary experience of understanding language. Gricean and neo-Gricean approaches haven't the resources to explain aspects of our experiences of meanings in terms of the involuntary perception of them as monadic properties of the results of language transactions. As a result, they must instead attribute to speaker-hearers participating in such transactions intricate patterns of communicative intentions and background mutual knowledge that have no echo in our experience of our own intentions, or of those intentions we naturally attribute to others when engaged in communication with them. Furthermore, the intentions that we do recognize ourselves or others to have when engaged in communication are either ignored or mischaracterized by Gricean and neo-Gricean approaches.

In chapter 9, I show how the folk picture of a common language that is induced by our experience of language and language understanding makes possible (and natural) practices of modifying the properties of ordinary-language sentences and terms that occur in the sciences and in mathematics. That is, our experience of language explains why it's so easy for us to think of artificial languages as extensions of ordinary language, differing at most in the introduction of vocabulary rather than as radically discontinuous with natural languages both semantically and syntactically. One crucial aspect is the invisibility to us of the rich structure of our own language capacities and, consequently, the invisibility of the rich syntactic and semantic structure of natural-language events. Another thing that's crucial is that we see natural languages as composed of public entities that can be with mutual agreement easily changed in their semantic properties.

In addition to these chapters, there are also two methodological interludes, one after chapter 4 and the second after chapter 6. I have distinctive views about how scientific methodology plays out in the language sciences—in particular (but not restricted to) the evidential role that speaker-hearer intuitions should have, as well as details about how special sciences operate. I've largely segregated these discussions in methodological interludes, where I can use the previously material to illustrate these more general points about methodology in a way that isn't possible at the very beginning of the book (in chapter 0). I also recommend, in this regard, the General Conclusion. It provides an over-all summary of my views, although in a way that presupposes the book has been read; nevertheless, an initial reading of it may be helpful.

One last point: you may find disconcerting an important aspect of philosophical exposition, as I practice it. Almost all philosophers—after characterizing one or another class of purported items as *not* existing—avoid discourse that informally quantifies over or otherwise refers to such items. Alternatively, if they *must* refer to such items or otherwise quantify over them, they endeavor to trap all such locutions within intentional verb forms, such

as “S believes that” There are exceptions, of course—brazen Meinongians for example—but this is standard expository practice nevertheless.

Discourse that violates this expository standard is indispensable, however (in particular, see Azzouni 2010a). Quantification over mathematical entities indispensably infuses scientific discourse, for example. And so, the denial of the existence of abstracta must accompany ontologically neutral quantification over them. A crucial corollary of this indispensability is that our *explanations* of phenomena must involve the same indispensable language that occurs generally. So, for example, references to mathematical entities must occur in scientific explanations despite the frank avowal that these don’t exist.

Philosophical prose is no exception to this. (How could it be?) You will see, therefore, the discussion of entities accompanied by, in some cases, the denial of their existence. Many of you will not share my numerous ontological aversions, but this book will be profitable to you despite this because of the other substantial theses about language use that I argue for. In any case, or so I claim, there is no recasting the explanatory discourse of philosophy to eliminate quantification over, and reference to, nonexistent entities.⁶

6. I suppose I should add—because many philosophers continue to be confused about this—that no Meinongian doctrines are being embraced in this book or by me. The entities I deny exist I deny the machinations of in every sense, metaphysically speaking. They have no properties; they don’t participate in truth makers; they are truly nothing at all. Nevertheless, there are truths about them, and I hope, there are some truths about them in this book. More accurately, there are truths in this book that involve nonreferring terms and that quantify over nothing at all. They are, I hope, truths nevertheless.

Methodological Preliminaries

0.1 Type–Token Confusions

Russell (1912, 93–4) once wrote:

Seeing that nearly all the words to be found in the dictionary stand for universals, it is strange that hardly anybody except students of philosophy ever realizes that there are such entities as universals.

Russell’s quick and self-congratulatory way of establishing the existence of universals—by an ontologically weighty use of the “stands for” or “denotes” relation—has been challenged in its specifics. It remains true, nevertheless, that most philosophers are committed to one or another kind of universal and that this is false of almost everybody else.

Actually, the situation is even stranger than the previous paragraph indicates. If “universal” is understood broadly—as covering abstract objects of all sorts—then interactions with universals seem widespread. Consider this remarkable observation once made by George Boolos (1998, 128):

It would be a rather demented philosopher who would think, “Strictly speaking you can’t see *The Globe* [a Boston newspaper]. You can’t even see an issue of *The Globe*. All you can really see, really immediately perceive, is a copy of some issue of some morning’s *Globe*.” To say this, however, reflects a misunderstanding of our word “see”: more than a misunderstanding, really,

it's a kind of lunacy to think that sound scientific philosophy demands that we think that we see ink-tracks but not words, i.e. word-types.

But that's hardly the least of what the puritanical nominalist will regard as illicit trafficking with the nonexistent. Boolos notes a page later that we "deal with abstract objects *all the time*." We listen to *radio programs*, we write *reviews of books*, we correct *mistakes*, and some of us draw *lines* in the sand and write *numbers* on chalkboards. To think otherwise about numbers on chalkboards, Boolos suggests, is to be confused about the word "on."

What's notable about all this is that, despite the apparent obviousness of Boolos's remarks, it's still true—as Russell claimed—that it's largely unnatural for nonphilosophers to recognize themselves as ontologically committed to universals or even as ever *talking about* them. Many philosophers have forgotten that many people—even educated ones—have never even *heard* of universals as such. And it's not just that, when they are described, people fail to respond with, "Oh right—*them*." Many philosophers (that is) have forgotten how foreign and odd the notion of a universal is despite its venerable and ancient pedigree. They've forgotten that people—even intelligent and well-educated ones—have trouble understanding the idea of a universal, and that they have to be prodded to see what's puzzling about such objects and why being concerned with them is important. Of course, *everyone*, nearly enough, has *heard* of numbers. But, as I note momentarily, if people are successfully pushed by philosophical dialogue into classifying numbers as *some kind of thing* they usually describe them as *ideas*.

So, start with the type-token distinction that Boolos makes so much of. This distinction emerges strikingly late in the philosophical tradition—with Peirce (1906), if I'm not mistaken. In any case, it must be introduced with some care to those unfamiliar with philosophy. Surprisingly, though, it's not *types* but *tokens* that nonphilosophers must be brought to recognize and attend to. Sentences—as types—for example, seem easily referred to in practice, and their properties seem easily indicated by speakers; it's the tokens, individual utterances (as sounds, chalk marks, carvings in stone) that we must work to get them to appreciate the properties of and to distinguish from the types that the nonprofessional so easily confuses them with. For example, the *sentence* has five words; the *sentence token* is made of chalk. The *sentence* begins with the word "the"; the *sentence token* has been uttered rapidly. Sentence-tokens come and go; sentences don't. Quine used to make a big deal of use-mention errors and their ubiquitousness. Type-token confusions are probably more widespread, and, indeed, such confusions seem—as it were—*enshrined* in the very grammar of natural languages, as I'll soon illustrate.

But first, consider the numeral-number distinction that Boolos also alludes to. Here, too, we have a distinction that it takes some work for the nonphilosopher to become nimble with. When the child—this is my recollection, anyway—is introduced to the distinction early in elementary school,

it's the notion not of *number* that the child needs to work to grasp but of a *numeral*: "5 is smaller than 10," but the numeral 10 is smaller than the numeral 5.¹ A numeral, we learn in elementary school, can be colored, but a number can't. I suspect that for the nonphilosopher—at least in this case—the number–numeral instance distinction is understood along the same lines that the type–token distinction is. A subsequent sophisticated viewpoint, however, treats these as logically heterogeneous: it assimilates the former distinction to one of (uses of) names or descriptions versus what's designated and the latter to instantiations versus types instantiated. In both cases what's striking is that ordinary practice—what we are naturally aware of and how we naturally describe what we are aware of—involves an easy "grasp" of certain kinds of universals, but it's only with a difficult shift in focus that we grasp their instances.

You may be puzzled: I started off two paragraphs ago with the claim that it's relatively unnatural for nonphilosophers to recognize themselves as committed to universals. Yet I've just finished claiming that in practice, at least in certain cases, nonphilosophers must be taught to focus on *instances* as opposed to the universals they so naturally refer to and describe the properties of. That's right, but when nonphilosophers reflect broadly on what they take there to be, they will express—ontologically speaking—"concreta prejudices." They will naturally *take themselves to see*, and naturally describe what they see as, specific objects located in specific places: a specific newspaper, specific sounds from a specific radio, specific chalk marks, and so on. It will take some work to make them understand that "sentence"—*as they use the term*—applies to types and not to tokens and that this contrasts with terms like "tiger" or "apple," which refer to individuals and not to types. "Maybe so," they will think—once they have gotten onto the claim just made—"but this is just a manner of speaking, and of course all I see and really mean to be talking about are those chalky things on the blackboard." Similarly, once we have pressed them enough on numbers so that they realize that, whatever it is that they're talking about when they do addition, it isn't things on paper, they will still refuse to desert their ("naïve") presupposed ontology by instead describing numbers as *ideas* or *concepts*. And they will *not* mean by "concept" Fregean concepts—the latter are abstract

1. This is still wrong, although it's significant that elementary-school teachers introduced the distinction to me as just described. *Numerals* are types *too*: what's smaller than what in the previous instance of the sentence phrase "the numeral 10 is smaller than the numeral 5" is a *token* of the numeral 10. I'm not *just* being pedantic by pointing this out; I'm noting that the *language* we use rather naturally drives us to focus on types, so we have to work hard—introduce philosophical jargon (e.g., "token" or "instance")—to even *talk about* the items Boolos's "rather demented philosopher" thinks are all that we see. Interestingly, the words we so introduce as jargon often have to be awkwardly extended in their usage to do the job needed. "Inscription," for example, can be drafted to cover concrete instances of sentences—but it's rather unnatural if extended to verbal utterances. "Utterance," has the converse problem. Schiffer (1987, 289) writes: "I shall continue Grice's artificially extended use of 'utter', which is intended to cover any behavior by which one means something." I'll try not to artificially extend it quite this far—but "utterance" and "utter," as I use them, will often include acts of writing.

objects. They will mean something that's vaguely psychological, something that—vaguely described—is “in the head.” I'm suggesting, therefore, that the targeted “demented philosopher” that Boolos ridicules takes his or her methodological starting point from the natural although apparently inconsistent thoughts of the nonphilosopher.

I mentioned that type–token confusions seem to be enshrined in the grammar of natural languages. It's more accurate to say that such confusions are aided and abetted in natural languages by two facts, both noted by linguists and philosophers. Identical-looking sentences, first, can nevertheless be quite different in their subject matters, referring with the same words, for example, to something like collections or to individuals in collections. Anaphora, second, allows unacknowledged shifts in the subject being talked about.

Illustrating the first point, it's not natural for the nonprofessional to notice the difference in topic between “Cats are animals,” and “Cats are happy,” on one hand, and “Cats are popular,” and “Cats are numerous,” on the other, or the difference in subject matter between “Cats are extinct,” and “Every cat is dead.”

To illustrate the second point, consider the following sentence:

- (1) That number, the one written on the blackboard by the teacher yesterday, is prime.

In this, and in other numerous grammatical utterances, we find slippage between *instances* of numerals and *numbers*. So, too, there is also slippage between tokens and types in:

- (2) That sentence, the one written on the blackboard by the teacher yesterday, has no adjectives in it.
- (3) The red book is too heavy, although it was favorably reviewed, and the blue one is boring, although everyone is reading it.²
- (4) That recording is a good one because the trumpet player is unusual, and despite the background hiss on the tape.³

2. I borrow (3) from Pietroski (2005, 278).

3. I don't mean to suggest the simple hypothesis that anaphoric “slippages” can be neatly catalogued either as token–type ones or as name-used–item-referred-to cases. On the contrary, the different kinds of “items” that can be simultaneously referred to are numerous, as Chomsky (2000a, 16) indicates, such as, “The bank was blown up after it raised the interest rates.” He writes, “words are interpreted in terms of such factors as material constitution, design, intended and characteristic use, institutional role, and so on. Things are identified and assigned to categories in terms of such properties” (15). Thus, many such factors can be included and excluded simultaneously in these kinds of examples; items differing in their properties along many dimensions can be simultaneously referred to. Some other examples of interest can be found in Koslicki (1999, 449, 447): “Gold, which is valuable, was discovered by the Sumerians,” and “Potatoes, which were introduced into Ireland in the seventeenth century, contain vitamin C.” Also, we can say this: “Dinosaurs are extinct, but in their heyday they were animals to be reckoned with.” Here the slippage isn't between types of things and individual things (the type “gold” isn't being asserted to be discovered by Sumerians, nor is the type “dinosaur” extinct).

In some cases, where the token of a described type is an otherwise stable object—perceivable, say, independently of a classification under the type that it’s otherwise assimilated to (such as “book” or “dress”)—one is relatively clear about the distinction between type and token. The specific book being indicated is colored and heavy—but *that* wasn’t the item that was favorably reviewed. Strictly speaking, no *item* was favorably reviewed. But in cases where tokens are more elusive—because they are events or *worse*, for example, the electronic books discussed below—it’s harder to attend to the tokens, to distinguish them from the types they are otherwise assimilated to, and to describe *their* properties. Consider a symphony. It takes a bit of puzzling to figure out what the *tokens* of a symphony might be: The score? Live events with musicians? CDs? CDs being played on a stereo? All of these? Nevertheless, symphonies seem relatively unproblematic to nonphilosophers: they easily discuss the properties of symphonies, and—without concern—confound reference to a symphony with references to, and descriptions of the properties of, whatever convenient tokens of such are at hand. For the philosopher, instead, various metaphysical worries lurk in the background about the identity of the various tokens of the types spoken of. There are the metaphysical worries in the family of the statue/clay sort, for example: the statue-shaped clay seems identical to the statue made of clay, yet one can survive the demise of the other.

It’s striking how little (not at all, I think it’s safe to say) nonphilosophers worry about identity issues—despite their ubiquitousness. The reason that they don’t seems to be evident, though. Nonphilosophers focus on tokens—insofar as they do focus on tokens—via the type classifications that those tokens are described in terms of. When focused on a token through its being a “statue,” the nonphilosopher thinks of it one way, and when focused on it through its being “clay,” the nonphilosopher thinks of it in another way. And, of course, the token can be focused on—in the same sentence—in *both* ways. Chomsky (2000b, 36) invokes the terms “filters” and “lenses” that provide “ways of looking at things and thinking about the products of our minds.” The image is a helpful one.

The slippage between the implicitly understood conditions on something being a token of *A* and the conditions on its being a token of *B* is overlooked because for the nonphilosopher (and for all of us when we’re simply engaging in conversation and not being “professional thinkers”) the type classification—psychologically speaking—compartmentally dominates the token even to the extent of partially dictating its “individuation conditions,” how far, that is, it extends in space and time. By “compartmentally” I mean that we evince no awareness of how the differing individuation conditions that different type classifications impose on tokens affect the metaphysical consistency of what we take ourselves to be talking about. We switch from one type classification to another effortlessly, even within the same “thought.” The *phenomenological compartmentalization* exhibited by our experience in this case is an extremely important systematic and widespread property of our conscious lives—of what we are aware of during a psychological episode and, equally important,

what we fail to be aware of during that episode. I'll exhibit the effects of phenomenological compartmentalization throughout this book.

Another indication of how it is that awareness of types apparently dominates our awareness of tokens, and not the other way around, is that the class of the paradigmatic tokens of a type—even when containing perceptually salient items that we do refer to—is allowed to mutate in ways that nonphilosophers find unproblematic. Consider:

- (5) I just downloaded the book that everyone has been talking about; it will be fun to read.

Electronic books are the unproblematic new tokens of the type *book* that have appeared in the last decade or so. It's not easy, however, to say exactly what these new tokens *are*—they are more event-like in certain respects than the traditional book. Certainly they exhibit a certain amount of ominous “metaphysical scatter.” But this, however puzzling to philosophers it may be, is rather unpuzzling to everyone else.

Even a little thinking about cases like these shows that, as types seem to arise in natural languages and as they are utilized and referred to by nonphilosophers, they aren't simple collections or sets of tokens that too steady a philosophical diet on examples like “tiger” and “electron” might suggest. Further, the natural picture of types as deriving their properties from the properties that their tokens all share (apart from those properties that types have by virtue of their being types and not tokens) is too simple.

My point is *not* to introduce a book on the topic of these puzzles—as philosophical puzzles to be studied in their own right—but rather to indicate in an introductory way a striking aspect of our understanding of language. It is—as I've indicated—as true of language types and tokens as it is of other types and tokens that we find it difficult and unnatural to distinguish them cleanly. To the extent that this happens, the nonphilosopher seems to have a firm grip on types but finds the tokens—that he is otherwise prone to claim are all there are—elusive of sharp characterization.

Because of this, language events and objects (for example, when we hear or see “sentences” and “grasp” what they mean) are far more like the elusive events that token *symphony* than they are like the traditional familiar squat objects that token *book*. So the phenomenology of language perception—what we are capable of recognizing ourselves as conscious of, and speaking about, in this domain—seems replete with the “awareness” of universals and their properties. This way of talking about sentences and their parts—and the accompanying phenomenology—echoes, in turn, the subpersonal processing of language, at least at a certain level of description. One possibility, thus, is that our subpersonal processing of language is replete as well with the subpersonal “grasping” of types—or at least the utilization of some sort of “prototypes”—that function as representations of various sorts.

Many contemporary theories of language, of course, simply take for granted that natural languages—whether they understand these as common languages or individualistically—are composed of types and, therefore at least at an important level of description, take us as engaged with such types. This is true even of the many theories that handle the contextual sensitivity of language via the role of the tokens of expressions—for example, those sensitive to the role of tokens of various words, such as “I.” Such tokens play a role in semantics via the “use” of expressions in “contexts” according to (for example) Kaplan’s very influential terminology. What enables them to do this, their “character,” however, “applies only to words and phrases as types” (Kaplan 1989a, 524). Note also that the popular terminology “use of an expression” treats tokens as akin to passing events or perhaps as akin to the temporal *stages* of the sentences they are the tokens of—like the “use of a screwdriver.” Conversely, that same phrase “use of an *expression*” seems to treat the type as the abiding *object*. The significance of this point will become clear in section 2.1.

Philosophers, when they become aware of the complications of type–token confusions and more generally of the individuation problems that words seem to present, become focused on solving these problems. They recognize them to be issues of *ontology*.⁴ Furthermore, they often describe the nonphilosopher as “confused” about certain distinctions, ones that it takes a philosophical slant to become clear about. This is decidedly not the approach of this book. That nonphilosophers are confused about certain distinctions (but aren’t confused about others) is evidence about their experience. In particular, that nonphilosophers are so utterly unaware of the type–token distinction in their thinking and usage that they stumble conceptually when they attempt awareness of it is here understood to be significant evidence about our conscious experience of language events—something I’ve indicated in the course of this section. It’s clear already that the phenomenology *collapses* the distinction in curious and systematic ways; it isn’t simply that the nonphilosopher has trouble keeping the distinction in mind. This calls, therefore, for diagnosis, for an explanation of what it is about our experience of language that makes the type–token distinction so elusive to us.

I’ve thus illustrated the diagnostic orientation of this book in this section with respect to one philosophical puzzle case. There will be others.

0.2 Indispensability and Ontological Commitment

By discussing, in section 0.1, how universals and types seem to intrude into our thinking about things, I am *not* repeating—in a new context—the old observation that we are particularly good at recognizing similarities and grouping things into kinds on the basis of those observed similarities. To notice that

4. See, for example, Epstein (2009) and the references therein.

we naturally group animals into kinds based on certain similarities salient to us—and even to suggest that such tendencies are innate—is nevertheless not necessarily to presuppose the existence of universals of any sort. (Arguments like *that*, when offered by philosophers, are embarrassingly slick.) To go from perceived similarities among things to the existence of a universal that such things share requires additional philosophical considerations—and such a move, as we’ve seen, *doesn’t* naturally occur to nonphilosophers. Indeed, on my reading, anyway, Plato⁵ explicitly introduced universals as explanatory tools, and some of the things he meant to explain such as what tigers have in common are still offered by contemporary metaphysicians as reasons to believe in universals. Such purported explanatory needs are dubiously motivated, but, apart from that, they are certainly artificial in the sense that a sophisticated, if not sophistical, notion of explanation must be employed. It doesn’t occur to the nonphilosopher that the fact that things are alike or that they can be grouped together—books, tigers, electrons—itself *requires* an explanation. This shows, if not that “explanation” is being stretched illicitly, at least that the explanatory puzzle that universals, in this context, are invoked to solve doesn’t seem pressing without substantial philosophical preparation.⁶ I won’t be writing much in this book about this particular route to abstract objects.

Ontological commitment to universals is motivated by another consideration that’s much harder to dismiss. This is the traditional and long-standing assumption that truths require what they are about to have properties so that those truths are “made true” by their rightly ascribing properties to those objects. Geometry—a previously empirical subject—was discovered to be deductively tractable once its concepts were massaged into the right form (for example, no-dimensional points, one-dimensional straight lines). The resulting concepts, however, have no “real-world” instantiations. Thus, this is, perhaps, the earliest example of the use of a class of valuable empirically applicable true sentences to force ontological commitments to what such sentences are “about” (see Azzouni 2004a).

I consequently trace a version of the “Quine-Putnam” indispensability argument back to ancient Greek times. This indispensability “argument” is actually a family of strategies that, on the basis of a commitment to mathematical truths that are indispensable to empirical applications, attempts to force a commitment to the abstract objects apparently referred to by terms in those truths. Aristotle, arguably, responded to Platonic versions of these arguments by rooting the truth of mathematical statements derivatively in truths about ordinary objects. Thus, a mathematical truth has as its relata—not universals but—actual space-time objects that instantiate

5. I’m thinking specifically of *Phaedo* (96a–102b) (Plato 1961a, 78–83).

6. It’s something in the neighborhood of such a bogus claim about the explanatory need for abstract objects that Quine (1953a, 10) dismisses when he writes: “That the houses and roses and sunsets are all of them red may be taken as ultimate and irreducible, and it may be held that McX is no better off, in point of real explanatory power, for all the occult entities which he posits under such names as ‘redness.’”

the properties, numerical or geometrical, that are focused on in mathematical contexts to the exclusion of other properties of those objects. (I draw this interpretation of Aristotle from Lear 1982.) I won't dwell further on this kind of response to the ancient form of the indispensability argument except to say that it faces fatal problems with respect to contemporary cases of applied mathematics (because instantiations of mathematical properties—in objects that we are otherwise ontologically committed to—aren't available).⁷

Let's notice instead that two families of strategy can be used to force an ontological commitment to abstract objects on the basis of statements held to be true. The first relies on Quine's criterion for ontological commitment. It therefore makes ontologically salient the "existential consequences" of a true theory:⁸ a useful piece of jargon that I'll adopt is to describe the ontological commitments of a piece of discourse as what its "quantifiers range over."⁹ The second strategy attempts to force commitments by means of an analysis of the notion of truth. This can be done directly by showing that the cogency of truth attributions to a certain class of sentences relies on a presumption of a correspondence of some sort. Thus a commitment to universals arises via a more general correspondence to objects, facts—or whatever—that are the truthmakers for truths. One can also proceed more indirectly by showing that the semantics of truth attributions presupposes, in one way or another—and among other things—universals.

With respect to the first strategy, an obsession with how ontological commitments to universals arise from applied mathematics—especially in physics—has given many philosophers of mathematics the impression that a piecemeal response to the indispensable applications of different branches of mathematics may free us of such commitments. This response to unwanted ontological commitments, at least in its methodology, is easy to understand and is due originally to Quine.¹⁰ Quine attempted to avoid ontological commitments to abstract objects, when possible, by showing that the locutions that commit us to such objects are—sometimes contrary to appearances—dispensable. "Paraphrase" is the name of the activity of replacing such objectionable locutions with others that can do the job needed but don't have undesirable commitments. Despite its roots in ordinary-language practices,

7. Aristotle, it may be thought, was unduly generous to himself by positing kinds of nonmathematical objects that could play the needed instantiation role vis-à-vis abstracta properties. Analogous accusations can be directed at contemporary nominalist programs (e.g., Field 1980) that attempt similar maneuvers with space-time points.

8. The true theories under discussion are usually amalgams of applied mathematics and an empirical subject matter—for example (well, "i.e.," pretty much among philosophers), a branch of physics (see Azzouni 2004b and references therein).

9. This jargon originally characterized the "ontological commitments" of an interpreted first-order language—when Quine's criterion is applied to it—because that criterion focuses on the existential quantifier. Use of the jargon is extended to natural languages by understanding intuitively recognized ontological commitments as ones that would—upon regimentation into a first-order formalism—be identified as such by Quine's criterion.

10. See the references under "Paraphrase"—and elsewhere—in the index to Quine (1960a, 291).

“paraphrasing” has a technical appearance in the hands of Quine—and other philosophers he has influenced. This is because of a widespread practice of replacing verbal locutions in the vernacular with ones from one or another interpreted formal language.

It’s natural, especially in the wake of Chomsky’s often highly technical work, to think of linguistics as like physics at least insofar as formal tools—in particular, the mathematics of formal grammars—are to be applied in the study of language. The analogy is an exact one: any “physical theory” can be treated as a branch of pure mathematics if it’s sufficiently mathematized and if it’s treated independently of empirical considerations—as a purely “deductive science,” as it were (Azzouni 1994, 108–9). One is engaged in physics, however—and *not* mathematics—if one doesn’t rest with the deduction of various results but is instead concerned with the empirical application or testing of such results (with their “empirical adequacy,” as it’s sometimes put). So too, formal grammars—however specified—are, in one notable tradition of linguistics stemming directly from Chomsky’s work, of interest only insofar as the study of them is included in a broader concern with the degree to which such grammars are “psychologically real.” What’s meant by describing such grammars as “psychologically real” is that “they constitute accurate models of the (implicit) knowledge that underlies the actual production and interpretation of utterances by native speakers” (Chierchia and McConnell-Ginet 1990, 2).¹¹

The suggested analogy between physics and linguistics hints that the problems facing the opponent of universals in linguistics are the same as those facing the opponent of universals in physics. Indeed, the universals that arise in the analysis of formal language—both the universal-stuff that such formal languages are themselves defined in terms of and what their semantics require there to be (what their terms refer to)—seem replicable in set theory. So it can be thought that any reasons for thinking that paraphrase will enable the elimination of universals (that is, set-theoretic constructions) from physicalistic discourse apply equally well to whatever universals (that is, set-theoretic constructions) arise in the context of language. I won’t be dwelling on the large industry, as that subject has arisen in contemporary philosophy, engaged in the paraphrasing away of quantification over abstract objects in the various empirical fields that mathematics is applied to. I’ve already done so elsewhere (see, e.g., Azzouni 2004b, 2009a).

Instead I’ll briefly summarize my views about this as follows: assuming a criterion for what a discourse is committed to, that is, Quine’s criterion—as such approaches invariably do—I’m sceptical that paraphrase strategies can really succeed. Two problems invariably arise: either the purported paraphrases help themselves to additional logical resources that make it hard to

11. Of course, “modeling” here—as so often in the sciences—involves a subtle interplay with various “idealizations” and “tentative working hypotheses.” For some of those operative in this case, see Chomsky (1986, especially 16–17, 36–37) and also Chomsky (1995, 6–7).

evaluate their purported commitments;¹² or a commitment to universals still remains, although an attempt is made to disguise this.¹³

I've argued (Azzouni 2004b) that nominalism doesn't require the elimination of quantification over universals in order to avoid ontological commitments to them. The key move is instead to reject Quine's criterion for what a discourse is committed to and to replace that criterion with a characterization of what ontological commitment actually comes to. In general, therefore, the indispensability of a referential idiom is insufficient to establish commitment; what's required for the latter is a (collective) attitude toward the purported reification in question. If the relata of a nounphrase, for example, are treated as things independent of us, in the sense that we have an epistemic practice in place of discovering their properties—an *epistemic practice that is sensitive to the items studied*—then those items are taken by us to be the sorts of things open to existence claims.¹⁴ Thus, for example, ordinary macro-objects, such as furniture, large animals, and various kinds of plant life, are taken by us not only to be susceptible to existence claims but also to actually have been shown to exist. The same is true of many of the unobservable posits of the empirical sciences, although subtler considerations are needed to establish this.

On the other hand (to choose an easy case), fictional characters aren't so taken by us to exist. A symptom of this is that, when such characters are being invented for the first time, their inventor can stipulate their properties without fear of falsification. Despite the absence of stipulation, the same is true of dream figures and hallucinations. Although, in these cases, we cannot dictate the properties of such objects, our collective assumption is that truths about them aren't dependent on actual objects that we are endeavoring to correctly describe the properties of.

The universals of mathematics are a much harder case. Due in part both to the rigidity of the proof procedures of mathematics and to the fact that mathematics is couched in the idiom of reference to objects, the objectivity of proven theorems is accompanied by the *prima facie* comprehending of objects that such theorems seem to establish the existence and the properties of. Nothing in the practice of mathematics, however, supports the assumption that *epistemic* processes involving sensitivity to objects can be found anywhere in that practice (see Azzouni 2004b, chapters 6 and 7). As I've said, I don't intend to revisit the arguments for this here. I'll take it more or less as a given both that universals don't exist and that—despite this—quantification

12. The many approaches that help themselves to one or another modal idiom fall into this category, as do approaches that use higher-order logics.

13. Field (1980)—despite a purported nominalism and despite the employment of a higher-order logic—in *addition* remains committed to universals (e.g., his already mentioned space–time points) in the sense that such are nevertheless quantified over (as many have pointed out). See Azzouni (2009a) for a detailed discussion of Field's program and its drawbacks.

14. "Exists," as used here and in the rest of this book, *usually* indicates an ontological concern. That the word is *often* not so used in the vernacular is something most readers, one hopes, don't need pointed out. See Azzouni (2010b).

over them is ontologically innocuous and thus doesn't require the surgery of paraphrase.

To this extent, therefore, the implicit viewpoint of the nonphilosopher is vindicated: reference to (and quantification over) abstract objects is consistent with a disbelief in them. Ontological commitments among us are *not* to be recognized by how one speaks—not even by how one speaks about what one “sees.” They are to be recognized by what one takes oneself to be *really* interacting with rather than by what one *says* one is interacting with. The latter is far too influenced by the demands of language—the intrinsic structure of such, the constraints on successful communication, ease of locution, and so on—to be trusted. Despite the richness in texture of what's quantified over, there really just are causally active objects. These are the things that we think really *exist*. Such things impact palpably on what's around them (including people). More basically, they are independent of us, and we need intelligent commerce *with them* in order to determine what their properties and powers are. The rest of what we “refer to” is—as it were—just “talk.”

The upshot is this: I'll be talking freely throughout this book of objects of various sorts without any commitments to such things existing. I'll also allude to various regularities about our experience of language transactions, the sharp formulation of which will require quantification over such objects. This is because quantification over such objects is central to our experience of language transactions themselves and because it's also crucial to language-processing theories that are currently the best candidates for explaining the psychological mechanisms leading to our experience of language. Sentence types are among such objects, but other sorts of objects with semantic properties are involved as well. I'll say a little more about this at the end of section 0.4.

0.3 The Relationship of the Science of Semantics to the Science of Psychology

The indispensability of quantification over universals is only a special case of a problem that arises in the sciences generally: that the terminology of a special science is often irreducible to that of the sciences “below it.” “Below” is metaphor for the metaphysical perspective that the objects and events characterized and studied in the special science in question come to (this is *still* metaphor, I'm afraid) “nothing more” than the objects and events characterized in the science below it. Special-science terminology often exists for the same reasons that (applied) mathematics has terms that refer to universals. This is to provide deductive tractability; and to make possible generalizations about the patterns and uniformities of behavior of phenomena that are otherwise invisible or inexpressible. An additional reason often present in the special sciences is that the particular form that evidence can take for the generalizations of the patterns and uniformities of the behavior of phenomena aren't

applicable to characterizations of entities in the underlying science. In other words, the special science terminology is evidentially “indispensable.”

Paraphrase, as understood in section 0.2, is the attempt to avoid apparent ontological commitments that arise because of indispensable terminology. Various reductionist programs—of one “science” to another—can share this motivation, although there are invariably other philosophical motivations as well. For example, various “physicalistic” doctrines are often motivated not merely by the view that what’s describable by physics must be—in some sense—all there is but also by the broader and perhaps vaguer claim that a physics’-eye view, if sufficiently complete, should have the resources to “explain” in principle everything that happens. Reductionist views of more local sorts can be held for motives that are more local versions of the above. Recanati (1993, 20), for example, expressing a widely held view, writes:

I believe, with Grice [*Studies in the Way of Words*, 1989], that meaning properties are to be analysed ultimately in terms of psychological properties. Sentence meaning is to be analysed in terms of utterance meaning, utterance meaning in terms of communicative behaviour, and communicative behaviour in terms of intentions and other psychological states.

Here, the motive behind the requirement of such an “analysis” seems to be that it’s not words or sentences that do anything semantic but only the people using them. Davidson (1990) writes (and Chomsky 2000e, 136 quotes him): “We all talk so freely about language, or languages, that we tend to forget that there are no such things in the world; there are only people and their various written and acoustical products” (also in Davidson 1992, 108). Grice (1989b, 340) says, “What words mean is a matter of what people mean by them.”

Anything apparently going on with the semantics of types, where those types are public-language sentences, therefore is ultimately reducible—perhaps in the strongest sense possible according to some and perhaps only in a weaker sense of what explains what according to others—to the goings-on in the psychology (and, ultimately, the brain physiology) of individuals while they engage in speech acts. Any reduction of public-language types to speech-act tokens, however, occurs simultaneously with a reduction of the meaning properties of those types to intentions and other psychological states that are the sources of communicative behavior.¹⁵

15. Schiffer (1988b) offers two motives for “intention-based semantics” (see Schiffer, xi: a program (i) of explicating a notion of speaker meaning in terms of acting with the intention of affecting an audience in a certain way and requiring that (ii) the semantic features of marks and sounds—“expression meaning”—can be explicated without relying on any other semantic notion apart from speaker meaning). The first motive is that “the sequence of marks ‘La neige est blanche’... has meaning among certain people, and that contingent fact... can hardly be regarded as a brute, primitive, and irreducible fact. Surely, the sequence has its semantic features by virtue of the communicative practices that prevail in the population for whom it has those features. But if meaning supervenes on use, then it’s reasonable to suppose that the use on which it supervenes can be described otherwise than in terms of the meaning supervening on it. Now it is further plausible that the use

Of course, the oft-repeated slogan to the effect that “words don’t mean anything; only people using words mean things” is compatible with the view Dummett (1986, 473) urges: that “words have meanings in themselves, independently of speakers.” Nevertheless (although I’ve not counted heads recently), it strikes me that the view Dummett urges is the minority one:¹⁶ most practitioners deny the existence of “common languages,” such as English and French, except in the sense that the specific languages of speakers overlap sufficiently that those speakers can take themselves at times to be among “fellow-speakers” and that therefore talk of such “languages” can be taken as an at times harmless “idealization.”¹⁷

A view similar both in its ultimate aim of couching semantics in psychology (or, ultimately, in neurophysiology) and in its method of individuating languages by individual speakers (although not similar in the actual pathway to the goal urged) is attributed to Chomsky by Pietroski (2005, 255–256):

Chomsky offers a plausible though often ignored conception of linguistic meaning and its relation to truth: the meaning of a natural language sentence *S* is an internalistic property of *S*, determined by the human language faculty and the relevant lexical items; the semantic properties of sentences, which reflect how human beings understand natural language, are theoretically tractable; but if an utterance of *S* is true or false, its truth or falsity is typically a massive interaction effect due to the meaning of *S* and many factors not indicated by elements of *S*.

Meaning, as Pietroski understands it, should be detached not only from truth but also from reference; nothing of *that* sort can be intended by Recanati. But both agree that language is to be rooted ultimately in psychology—personal and subpersonal. On the Chomskian view,¹⁸ the apparatus associated with the notion of a common language (for example, conventions and regularities of language use) are theoretically intractable and should be avoided altogether. Theorists like Recanati (1993), however, seem open to the interpretation that they treat these notions as convenient but temporary stopping points on the yellow brick road to an ultimate analysis.

There is a sense in which the sort of nominalist I aspire to be cannot disagree—on sheer grounds of nominalism, anyway—with the denial of the existence of common languages. The sentence types that I’ve described

which leads to meaning has to do with conventional regularities in acts of speaker-meaning . . .” (xii). The second motive is a commitment to a reductive form of “physicalism”—that the truths couched in semantic and psychological terms can be so stated in “physicalistic or topic neutral terms” (ibid.).

16. And, despite Davidson’s (1986, 446) remark that “there is no such thing as a language, not if a language is anything like what many philosophers and linguists have supposed” and Chomsky’s (2000e, 136, italics in original) remark that “to most philosophers of language, it is . . . obvious that there *are* such things in the world as languages: indeed, ‘common, public languages’—Chinese, German, etc. . . .”

17. It’s striking that Dummett regards the concept of a “[common] language” to be “indispensable” (Dummett 1986, 465–466) and the individualistic alternative that Davidson (1986) is offering to be “absurd” (474). I try to make something of this shortly.

18. See, for example, Chomsky (1986, especially section 2.4) and (1993); see also Pietroski (2005, 266).

nonphilosophers as seemingly engaged with—in the way they speak about words and sentences and in their language-processing phenomenology—don't exist. What's *really* going on, therefore—all that can really be going on—is the communication of individuals with one another by way of verbal acts, along with various accompanying psychological states and brain-physiological processes. If an analysis of such language events in terms of truth conditions and items referred to is theoretically intractable, as Chomsky and Pietroski suggest, this motivates a model of language as the internal processing of lexical items. Otherwise, one can try to (ultimately) analyze the communicative situation in terms of the psychological states of the individuals, and the truth conditions of their speech acts. In neither case do common languages survive as a viable subject of study.¹⁹

The worry I have about attacks on the notion of a common language by proponents of “individualistic approach to language” lies somewhat but only somewhat apart from issues of nominalism. I share a concern related to Dummett's claim that there is an “indispensable notion” of a common language. My concern differs in that I claim there is an indispensable notion of *shared language*—an indispensable notion of collections of language entities and events that have meaning-properties perceived *in common* by groups of people. It's this indispensable notion that motivates the natural belief that people speak common languages, such as English or German, although the unavoidable impression that language entities and events are perceived by groups of people *doesn't imply* that these entities and events therefore belong to common languages like English or German. This book undertakes the extended argument that *indispensable* talk of *public-language* entities arises from the involuntary experience of such entities *whenever* we engage in language transactions; it arises from what can be called *fellow-speaker phenomenology*.

Dummett, like many philosophers sensitive to the social elements of language, alludes to how Putnam's “linguistic division of labor” and the normativity of usage—encapsulated in the idea that any one speaker has only a “partial, and partly erroneous, grasp of the language”²⁰—seem conceptually linked to the notion of a common language. Dummett, however, doesn't think of the common language, as so conceived, as composed of Platonic types. Rather, it is “an existing pattern of communicative speech” (Dummett 1986, 467).

This way of characterizing the notion of a language in common is troubled. However successful it turns out to be as a polemical tool against *Davidson's* particular individualistic replacements of the common language (Dummett's official target), it doesn't seem that this view has much in the way of resources to defend itself against *Chomskian* objections. The central

19. Of course, quantification over abstract objects doesn't disappear (the indispensability of such quantification guarantees that); it just takes other forms. For Chomsky (2000b, 26), for example, “a linguistic expression... is a complex of phonetic, semantic, and other properties.” See Bromberger (1992, 174–5) for a compelling argument that “linguistics must be concerned with facts about types.”

20. Dummett (1986, 468). See also Hacking (1986).