The Blackwell Guide to the Philosophy of Language

Edited by

Michael Devitt and Richard Hanley

Blackwell Publishing

The Blackwell Guide to the

Philosophy of Language

Blackwell Philosophy Guides

Series Editor: Steven M. Cahn, City University of New York Graduate School

Written by an international assembly of distinguished philosophers, the *Blackwell Philosophy Guides* create a groundbreaking student resource - a complete critical survey of the central themes and issues of philosophy today. Focusing and advancing key arguments throughout, each essay incorporates essential background material serving to clarify the history and logic of the relevant topic. Accordingly, these volumes will be a valuable resource for a broad range of students and readers, including professional philosophers.

1 The Blackwell Guide to EPISTEMOLOGY
edited by John Greco and Ernest Sosa

2 The Blackwell Guide to ETHICAL THEORY
edited by Hugh LaFollette

3 The Blackwell Guide to the MODERN PHILOSOPHERS
edited by Steven M. Emmanuel

4 The Blackwell Guide to PHILOSOPHICAL LOGIC
edited by Lou Goble

5 The Blackwell Guide to SOCIAL AND POLITICAL PHILOSOPHY
edited by Robert L. Simon

6 The Blackwell Guide to BUSINESS ETHICS
edited by Norman E. Bowie

7 The Blackwell Guide to the PHILOSOPHY OF SCIENCE
edited by Peter Machamer and Michael Silberstein

8 The Blackwell Guide to METAPHYSICS
edited by Richard M. Gale

9 The Blackwell Guide to the PHILOSOPHY OF EDUCATION
edited by Nigel Blake, Paul Smeyers, Richard Smith, and Paul Standish

10 The Blackwell Guide to PHILOSOPHY OF MIND
edited by Stephen P. Stich and Ted A. Warfield

11 The Blackwell Guide to the PHILOSOPHY OF THE SOCIAL SCIENCES
edited by Stephen P. Turner and Paul A. Roth

12 The Blackwell Guide to CONTINENTAL PHILOSOPHY
edited by Robert C. Solomon and David Sherman

13 The Blackwell Guide to ANCIENT PHILOSOPHY
edited by Christopher Shields

14 The Blackwell Guide to the PHILOSOPHY OF COMPUTING AND INFORMATION
edited by Luciano Floridi

15 The Blackwell Guide to AESTHETICS
edited by Peter Kivy

16 The Blackwell Guide to AMERICAN PHILOSOPHY
edited by Armen T. Marsoobian and John Ryder

17 The Blackwell Guide to PHILOSOPHY OF RELIGION
edited by William E. Mann

18 The Blackwell Guide to the PHILOSOPHY OF LAW AND LEGAL THEORY
edited by Martin P. Golding and William A. Edmundson

19 The Blackwell Guide to the PHILOSOPHY OF LANGUAGE
edited by Michael Devitt and Richard Hanley

The Blackwell Guide to the Philosophy of Language

Edited by

Michael Devitt and Richard Hanley

Blackwell Publishing

BLACKWELL PUBLISHING
350 Main Street, Malden, MA 02148-5020, USA
9600 Garsington Road, Oxford OX4 2DQ, UK
550 Swanston Street, Carlton, Victoria 3053, Australia

First published 2006 by Blackwell Publishing Ltd

1 2006

Library of Congress Cataloging-in-Publication Data

The Blackwell guide to the philosophy of language /
edited by Michael Devitt and Richard Hanley.
p.cm.—(Blackwell philosophy guides ; 19)
Includes bibliographical references and index.
ISBN-13: 978-0-631-23141-7 (hardback : alk.paper)
ISBN-10: 0-631-23141-2 (hardback : alk.paper)
ISBN-13: 978-0-631-23142-4 (pbk. : alk. paper)
ISBN-10: 0-631-23142-0 (pbk. : alk.paper)
1. Language and languages—Philosophy. I. Devitt, Michael,
1938– II. Hanley, Richard. III. Series.
P107.B582006
401—dc22
2005028555

A catalogue record for this title is available from the British Library.

Set in 10/13 Galliard
by SPI Publisher Services, Pondicherry, India
Printed and bound in India
by Replika Press Pvt. Ltd

The publisher's policy is to use permanent paper from mills that operate a sustainable forestry policy, and which has been manufactured from pulp processed using acid-free and elementary chlorine-free practices. Furthermore, the publisher ensures that the text paper and cover board used have met acceptable environmental accreditation standards.

For further information on
Blackwell Publishing, visit our website:
www.blackwellpublishing.com

Contents

Contributors

Kent Bach is Professor of Philosophy at San Francisco State University.

Simon Blackburn is Professor of Philosophy at the University of Cambridge.

Max Cresswell is Professor of Philosophy at the University of Auckland and Texas A&M University.

Martin Davies is Professor of Philosophy in the Research School of Social Sciences, Australian National University.

Jerry Fodor is State of New Jersey Professor of Philosophy at Rutgers University.

James Higginbotham is Linda Hilf Chair in Philosophy, and Professor of Philosophy and Linguistics, at the University of Southern California.

Paul Horwich is Professor of Philosophy at New York University.

Frank Jackson is Distinguished Professor of Philosophy and Director, Research School of Social Sciences, the Australian National University.

Ernie Lepore is Professor of Philosophy and Director of the Center for Cognitive Science at Rutgers University.

Brian Loar is Professor of Philosophy at Rutgers University.

Peter Ludlow is Professor of Philosophy and of Linguistics at the University of Michigan.

William G. Lycan is William Rand Kenan, Jr. Professor of Philosophy at the University of North Carolina.

Vann McGee is Professor of Philosophy at the Massachusetts Institute of Technology.

Alexander Miller is Professor of Philosophy at the University of Birmingham.

Stephen Neale is Professor of Philosophy at Rutgers University.

Karen Neander is Professor of Philosophy at the University of California, Davis.

John Perry is Henry Waldgrave Stuart Professor of Philosophy at Stanford University.

Mark Richard is Professor of Philosophy at Tufts University.

Stephen Schiffer is Professor of Philosophy at New York University.

Stephen P. Schwartz is Professor of Philosophy and Religion at Ithaca College.

Josef Stern is Professor of philosophy and of Jewish studies at the University of Chicago.

Acknowledgments

Editorial assistance was provided by Sophia Bishop, James Dow, and Scott Jones. Thanks to Timothy Czech for compiling the index. Richard Hanley's work on this volume was supported in part by the University of Delaware General Research Fund. Thanks also to Jeff Dean, Danielle Descoteaux, Graeme Leonard, and Nirit Simon.

Introduction

Michael Devitt and Richard Hanley

Foundational Issues

The philosophy of language is both fascinating and difficult. One reason for this is that hardly any issue in this area is uncontroversial. Controversy begins with some foundational and methodological questions. Consider, for example, this very basic question: What are the tasks of the philosophy of language? One obvious task is: the study of linguistic *meanings*. But this immediately raises two questions.

First, what are these "meanings"? Linguistic expressions have the function of communicating messages, conveying information about the world. Clearly, the meanings of expressions play a crucial role here. Yet, as Martin Davies notes in chapter 1, we cannot simply identify meanings with messages because the one sentence can be used to communicate different messages on different occasions; it can imply things that it does not literally say. We need to distinguish its literal meaning, studied by *semantics*, from other properties it may have that are studied by *pragmatics*. But there is much controversy about where and how to draw the line between semantics and pragmatics; see discussion in chapter 8.

Second, what sort of "study" do we have in mind? Is semantics empirical or is it a priori? Is it a science? In what way is it philosophical? These questions dominate Davies' discussion. Someone who supposes, as many do, that philosophy is entirely a priori, will think that semantic theorizing can go on independent of any science. This antireductionist view is what Davies nicely calls "philosophical isolationism." At the other extreme, naturalistically inclined philosophers will think that semantics reduces to empirical cognitive science. Davies calls this "cognitive scientism." He would like to find an intermediate position, as most philosophers probably would. But finding it is difficult. Paul Horwich discusses reductionism and antireductionism in chapter 2.

There is a related issue. The dominant method in semantics is to consult "intuitions" about what an expression means, refers to, and so on, intuitions

that are usually elicited in "thought experiments." What are we to make of this practice? The isolationist will think of this as the characteristic method of "armchair philosophy," yielding intuitions based on a priori knowledge of concepts. And she is likely to think that it is *the* task of semantics to account for these intuitions. The naturalist must see the intuitions as having the same empirical status that intuitions generally have in science and as serving at best as fallible evidence for a semantic theory.

Other foundational issues press in. What is the relation of language to *thought* (discussed in detail in chapter 4)? The folk idea that "language expresses thought" leads to the view, developed by Paul Grice (1989), that thought is explanatorily prior to language. Another influential view, developed by Donald Davidson (1984), is that thought and language are interdependent with the result that their explanations proceed together. In contrast to both these views, Michael Dummett (1993) gives priority to language. This approach is, as Davies remarks, "apt to sound rather behaviouristic." Even Davidson's approach starts from the behavioristic assumption that "meaning is entirely determined by observable behavior, even readily observable behavior" (1990: 314).

How does the obvious fact that competent speakers of a language "know," in some ordinary sense, the meanings of its expressions bear on our theory of the meanings known? Our answer to this will depend on what we make of the knowledge. It is mostly not explicit propositional knowledge but it is common to think of it, as Davies does, as "implicit" or "tacit" knowledge. This can lead to the view that the semantic task simply is the study of this state of knowledge: as Dummett puts it, "a theory of meaning is a theory of understanding" (1975: 99). Or it can lead to Davidson's view, noted by Davies, that a theory of meaning *suffices* for understanding. But perhaps we should think of the speaker's knowledge as mere knowhow, a cognitive *skill*, that need not involve any propositional knowledge, whether explicit or tacit. This view leads to a sharp distinction between the theory of meaning and the theory of the competent person's knowledge of meanings.

Davies sympathetically contrasts the Davidsonian and Gricean approaches to meaning. We note that each has been subject to criticisms of a foundational sort. Thus, the Gricean approach aims to explain linguistic meaning in terms of thought content but leaves the latter unexplained, as Brian Loar notes in chapter 4. This can seem unsatisfactory given that thought content is rather similar to linguistic meaning. Davidsonians have the idea that the semantic task is to spell out the conditions on axiomatic theories that correctly specify the meanings of expressions in particular languages, conditions on the "radical interpretation" of those languages. This interesting idea is rather taken for granted by Davidsonians and yet it is not obviously right. If we suppose that the task is to *explain* meanings in general, why suppose that we can accomplish this by studying the constraints on meaning *specifications* for languages?

For more on some of these foundational questions, see Devitt (1996: ch. 2).

Part II: Meaning

Theories of meaning

Moving from foundational to more substantive matters, we note the very important principle of *compositionality*: the meaning of a sentence is determined by the meanings of the words that constitute it and by the way those words are put together, by the syntactic structure of the sentence. This is a fairly uncontroversial fact about language, although, as Horwich points out (chapter 2), there is controversy about how to take account of it theoretically.

Compositionality enables us to explain an obvious fact about linguistic competence, as Davies notes. A competent speaker of a language is able to understand an indefinitely large number of sentences that are entirely novel to her. How? She understands the words in the sentences and has mastery of the syntactic rules that govern the structures of sentences.

An idea that seems to strike everyone when they first think about meaning is that the meaning of an expression is tied closely to the way in which we would *tell* whether the expression applies to something. The idea was captured in the logical positivists' slogan, "Meaning is method of verification." This verificationism was very popular in philosophy in the 1930s and 1940s, the heyday of positivism, but has since fallen from favor. However, a form of verificationism had a brief revival under the influence of Dummett (1993).

The most popular idea in philosophy for explaining meaning has been the idea that it is largely, if not entirely, a matter of explaining truth conditions. The idea is that the meaning of a sentence is to be explained by relating it to the circumstances under which it would be true, an explanation that will involve the referential relations of its words. Gottlob Frege, whose theory is described by James Higginbotham in chapter 3, is usually regarded as the father of the truth-referential approach to semantics. Bertrand Russell also had a truth-referential theory. There have been many different theories of this sort since, but an influential one in recent times has been Davidson's, inspired by Tarski's famous theory of truth (which is discussed in chapter 20). Davidson's basic idea is for specifying the meanings of the sentences of a language in that very language. We construct a theory with referential axioms like ' "Socrates" refers to "Socrates" ' and axioms for combining words into sentences, a theory which enables us to derive as theorems ' "*s*" is true if and only if "*p*"' whenever '*s*' is replaced by a canonical description of a sentence, and '*p*' by that very sentence. The basic idea then has to be extended to specifications of the meanings of one language in another language, and to cope with ambiguities and other complications, as Higginbotham brings out; see also Loar in chapter 4.

Compositionality leads us to expect that a truth-referential explanation of meaning will proceed in two steps. First, the truth conditions of sentences will

be explained in terms of their syntactic structures and the references of the words that fit into those structures. Second, the references of the words will be explained by theories of reference (of the sort discussed in part III of this book). So we expect theories of reference to be a central part of explaining meaning. A controversial consequence of Davidson's holistic interpretative approach is that there is no need for, nor possibility of, theories of reference.

There has been a problem from the beginning with supposing that the truth-referential properties of a sentence exhaust its meaning. For, as Frege pointed out, the meaning of 'Hesperus = Phosphorus' surely differs from the meaning of 'Hesperus = Hesperus.' Yet, according to the supposition, they should have the same meaning because 'Hesperus' and 'Phosphorus' have the same referent (the planet Venus) and hence the same meaning. This led Frege to introduce his famous notion of *sense*: although 'Hesperus' and 'Phosphorus' have the same referent – the planet Venus – they differ in sense or "mode of presentation," the former meaning, say, "the heavenly body seen in the evening" and the latter, say, "the heavenly body seen in the morning." This matter comes up often; see particularly chapters 10 and 14.

It is obvious that the meaning of an expression depends somehow on its use: if we had used the word 'cat' to refer to dogs instead of cats it would have had a different meaning. Inspired by this fact, and his "deflationist" view of truth and reference (see chapter 20), Horwich (1998) has proposed a "use theory" of meaning. The guiding idea is that the meaning of a word is engendered by its "basic acceptance property," the fact that specified sentences containing it are accepted underived. See section 10 of chapter 2 for a brief discussion.

Thoughts and meaning

We have already mentioned one issue that arises in thinking about the relation of thought to language. This is an issue of explanatory priority. Thus Grice thinks that thought content is explanatorily prior to linguistic meaning. Davidson, on the other hand, sees no priority, taking the two concepts to be coordinate. Davidson's view rests, as Loar notes in chapter 4, on "the principle of charity." According to this rather surprising principle, we should interpret another's sentences, and ascribe beliefs to another, so as to make the sentences and beliefs come out, so far as possible, *true*. The principle reflects the influence of Quine's meaning skepticism (see chapter 5) and Davidson's basically antirealist view of mind and meaning: meanings are not for the most part objective properties with natures awaiting our discovery; interpretation is more a matter of *imposing* a reality than discovering it. Loar discusses the views of Grice and Davidson, as well as the earlier views of Frege and Russell; see also Horwich in chapter 2, section 4.

As Loar points out, Russell's discussion (1989) of "internal speech" raises another issue about the relation of thought to language: Do we think in a language? And if we do, what language do we think in? These have been contro-

versial issues. Jerry Fodor (1975) has argued for "the language-of-thought hypothesis" according to which we think in a language-like system of representation. He rejects the idea, endorsed by Gilbert Harman (1973), that this system is the natural language of the thinker. For Fodor the language of thought is not the language of talk but rather an innate universal internal language, often called "Mentalese."

Meaning skepticism

Meaning skepticism is not simply the view that we cannot find out the facts about meanings. It is the view that there are no facts to be found out: there is no fact of the matter about what expressions mean. Clearly if it *were* a factual matter then it would also be a factual matter whether two expressions mean the same and hence whether one translates the other. So an argument that it is not a factual matter whether one expression correctly translates another is an argument for meaning skepticism. Quine's famous argument for the indeterminacy of translation is such an argument (1960). In fact, he had two arguments for indeterminacy, "the argument from below" and "the argument from above" (1970). Alex Miller discusses the two arguments and some responses to them in chapter 5. Whatever one makes of the arguments, it is important to note, as Miller does, that the terms of the debate are Quine's. Quine has a very restrictive behavioristic view of the sort of facts that *could* determine meaning.

The second argument for meaning skepticism discussed by Miller is one that Saul Kripke (1982) extracts from Wittgenstein. This argument has a much broader view of possible meaning-determining facts, allowing in mental facts. The argument is that no facts determine that a person using an expression is following one rule for its use rather than many others; hence that no facts determine that it means one thing rather than many others. Kripke's skeptic has a number of arguments for this view but the main one, briefly discussed by Miller, concerns normativity. The argument is that the dispositional facts alleged to determine the meaning of a term fail to do so because they do not tell us how we *ought to* apply the term. This argument has much exercised commentators.

The analytic–synthetic ("a–s") distinction

It has been common to believe that some true sentences are "analytic" in that they are true solely in virtue of meaning, whereas those that are not so true are "synthetic." Thus it was held that (U), 'All bachelors are unmarried', is analytic because 'bachelor' just means the same as 'unmarried man.' In contrast, (F), 'All bachelors are frustrated' is synthetic, depending for its truth on extralinguistic facts about bachelors. As Jerry Fodor and Ernie Lepore point out in chapter 6, the hope was that this would explain why a sentence like (U) is *necessary* whereas one like (F)

is contingent; and, on the assumption that we *know* the meanings of the terms in (U) and (F), why (U) is a priori whereas (F) is a posteriori. Aside from all this hoped for work, the a/s distinction seemed intuitively plausible. But there is a problem: surely, (U) is not true *solely* in virtue of meaning. The fact about the meaning of 'bachelor' shows that (U) is synonymous with 'All unmarried men are unmarried' but it is hard to see how that "logical truth" could be true solely in virtue of meaning. So how could (U) be? The definition of "analytic" had to be modified: a true sentence is analytic if it can be reduced to a logical truth by substituting synonyms for synonyms. (This modification dashes the hope of explaining apriority by analyticity: knowledge of the truth of (U) rests on knowledge of a logical truth and analyticity is no help in explaining that.)

Even the modified a/s distinction seemed to fail in the face of Quine's sweeping criticisms. Recently, however, Paul Boghossian (1996, 1997) has argued that someone who is a realist about meanings, hence the opposite of a meaning skeptic, must accept the a/s distinction. This argument is the main target of Fodor and Lepore. Their case starts from the premise that analyticity requires not simply concept synonymy but concept *identity*: for (U) to be analytic, the concept expressed by 'bachelor' must be identical to that expressed by 'unmarried male.' They argue that this is not the case: the former concept, BACHELOR, is simple, the latter, UNMARRIED MALE, is complex. They note that their starting premise might be challenged but they offer several considerations in its defense.

Formal semantics

Formal semantics uses the techniques of formal logic to throw light on the meanings of natural language sentences. The expressions of a formal language have a clear and transparent semantics: the relation between the syntactic forms of these expressions and the situations that would make them true is well understood. So if we can find a formal expression that "means the same as" a natural language sentence, then we can learn a lot about the meaning of the sentence. The formal paraphrase of a sentence is often called its "logical form." The key to finding these paraphrases is the earlier-mentioned principle of compositionality: the meaning of the whole is a function of the meaning of the parts and their mode of combination. This principle is central to formal semantics.

As Max Cresswell points out in chapter 7, compositionality can be nicely illustrated by definitions of 'not', 'or', and 'and'; for example, "not α is true if α is false and false if α is true." Things are not mostly that simple, of course. To take account of tensed sentences we need to consider truth at a time; to take account of modal sentences we need to consider truth at a possible world. In general we need to consider truth *at indices*. Cresswell demonstrates the success of these formal techniques in handling tensed sentences, quantifiers, and other features of language.

Speech acts and pragmatics

As already noted, we need to distinguish what a sentence means from what a speaker means when using that sentence. Thus, as Kent Bach points out in chapter 8, in a performative utterance one performs an act by uttering a sentence. For example, one can apologize by saying "I apologize." Some performative utterances are institution-bound conventions, such as a judge's "Overruled!" J.L. Austin identifies three distinct levels of action beyond the act of utterance itself. He distinguishes the act of saying something, what one does in saying it, and what one does by saying, and dubs these the "locutionary," the "illocutionary," and the "perlocutionary" act, respectively. An illocutionary act succeeds if the speaker's audience recognizes the speaker's intentions. As a perlocutionary act it succeeds only if the audience actually fulfills the speaker's request.

Grice's theory of conversational implicature aims to explain how a speaker can mean just what he says or can mean something more or something else entirely. His notion can be applied to illocutionary acts. When an utterance is performed indirectly, it is performed by way of performing some other one directly. When an utterance is nonliteral, what the words mean is not at all what the speaker means. However, Bach argues, Grice overlooks conversational "impliciture," where what the speaker means is implicitly conveyed rather than implicated, by way of expansion or completion.

Historically, the semantic–pragmatic distinction falls into three types: linguistic meaning vs. use, truth-conditional vs. non-truth-conditional meaning, and context independent vs. context dependence. Bach argues that the proper distinction can be drawn with respect to various things, such as ambiguities, implications, presuppositions, interpretations, knowledge, process, rules, and principles. The distinction applies fundamentally to types of information. Semantic information is information encoded in what is uttered together with any extralinguistic information that provides values to context-sensitive expressions in what is uttered. Pragmatic information is the information the hearer relies on to figure out what the speaker is communicating. This distinction is particularly useful in providing a simple account of how people can often communicate efficiently and effectively without the need to make explicit what they are trying to convey.

Figurative language

The oldest conception of metaphor characterizes it as improper or deviant use of the literal. This idea is undermined by the observation that some metaphors are equally true in the same contexts, whether interpreted literally or metaphorically. A second claim is that a sentence used metaphorically might have a different truth-value from what it would have were it interpreted literally. This entails that the same sentence must have a different meaning when used metaphorically than when

used literally. However, the meaning of a metaphor often cannot be understood without knowing the literal meaning of its utterance. This metaphorical-literal dependence is best understood by a theory of pragmatics. The literal can be identified as what a sentence 'S is P' means, whereas the metaphorical meaning is what a speaker can use it to mean, say, that S is R. Searle proposes that R is the metaphorical meaning of the predicate P on a particular occasion, and the fact that P conveys R according to some pragmatic principle is the sense in which the metaphorical depends on the literal. Josef Stern argues in chapter 9 that although Searle's account correctly demonstrates that there is no single ''ground'' that generates all metaphorical contents, it insufficiently explains why something is or is not a metaphorical meaning. Others, like Richard Rorty, claim that ''metaphor belongs exclusively to the domain of use.'' Similarly, Donald Davidson claims that ''a metaphor doesn't say anything beyond its literal meaning.'' But, Stern points out, if a sentence used metaphorically does not have a literal meaning its metaphorical effect cannot depend on it or be explained by its means; and if literal meaning is anything like truth-conditions, it is not at all clear that we know under what conditions many classical metaphors like 'Juliet is the sun' would be true.

Metaphorical meaning is context dependent. First, metaphorical interpretations of utterances of the same expression may vary widely from one occasion, or context, to another. Second, the interpretation of a metaphor is typically a function of all sorts of extralinguistic presuppositions, skills, and abilities such as the perception of similarity or salience. However, Stern argues, metaphorical interpretation does not come simply from looking at the content of each metaphor. Rather, only at a level that relates each content of the same expression used metaphorically to a relevant feature of its respective context of use, namely, the shared presuppositions, will metaphorical interpretations follow regularities. The presuppositions here are the sets of propositions to which a speaker, in making an utterance, commits himself, in the absence of which his assertion would be inappropriate or uninterpretable. Following this method, we see that a metaphorical interpretation or content is always fixed or constrained by its actual context of utterance.

Propositional attitude ascription

People often say things of the following sort: '*X* believes that *p*', '*X* said that *p*', '*X* hopes that *p*, '*X* wonders whether *p*', and so on. Since these seem to ascribe attitudes to the proposition *p* they are called ''propositional attitude ascriptions.'' These ascriptions are very important in our relations with *X* – they help us explain and predict *X*'s behavior – and with the world in general – if *X* believes that *p* and is reliable then probably p. The meaning of such an ascription is unusual in two interesting ways. First, it is not extensional (on one reading, at least). Thus, suppose that (T), 'Mary believes that Twain is witty', is true. Still (C), 'Mary believes that Clemens is witty', might be false. For, even though 'Clemens' and

'Twain' are coextensional – they both refer to the one person – Mary may not know this. Substituting a coextensional term in the 'that' clause is not guaranteed to preserve truth. Second, these ascriptions give rise to what Brentano called ''intentional inexistence.'' 'Octavia believes that Zeus destroyed Pompeii' might be true even though Zeus does not exist. These ascriptions can be true even though a term in the 'that' clause fails to refer. Attempts to explain propositional attitude ascriptions have played a major role in the philosophy of language.

As Mark Richard points out in chapter 10 the key issue in recent times has been how modes of presenting, ways of thinking about, an object enter into the truth conditions of attitude ascriptions. As we have already noted, Frege argued that a proper name like 'Hesperus' has a sense which is a mode of presenting the referent, say, ''the heavenly body seen in the evening.'' Russell thought that the meaning of a ''logically proper name'' would be its referent but thought that all ordinary proper names were truncated descriptions. So, in effect, his view of ordinary names was similar to Frege's. This view became known as the description theory of names. This theory provides 'Twain' and 'Clemens' with different modes of presentation which can then be used to explain the different truth values of (T) and (C). But, under the influence particularly of Kripke (1980), many came to think that the description theory was false: a name does not have a descriptive sense (see chapter 14). Where to go from there? One, sadly unpopular, response was to suppose that names have a nondescriptive *causal* sense or mode of presentation which can then be used to handle attitude ascriptions (Devitt 1981, 1996). A more popular response was that of ''direct reference'': there is no more to the meaning of a name than its referent. How then are we to explain the difference between (T) and (C)? Typically, appeal is made to the distinction between what (T) and (C) strictly say, which is alleged to be the same, and the other information that they convey, which is different. This is a distinction between semantics and pragmatics of the sort that we have just been discussing. But, as Richard argues, it is hard to make a solution along these lines persuasive. In particular, it is hard to make it compatible with the role of attitude ascriptions in commonsense psychological explanations. Richard's own explanation of attitude ascriptions is that ascribing an attitude to a person involves a sort of translation of the person's mental representation.

Conditionals

Conditionals are sentences of the form 'If A then B': in symbols (A → B). It is widely agreed that '~', '&' and 'v' ('not', 'and' and 'or', respectively) are truth functions: the truth values of a compound sentence formed using them is fully determined by the truth values of the component sentences. The simplest and oldest theory of the conditional holds that '→' is also a truth function, known as material implication, and in particular that (A → B) is equivalent to both of: (~A v B), and ~(A & ~B). This implies that the falsity of A and the truth of B are

separately sufficient for the truth of (A → B). Many find these results very implausible and they are known as the paradoxes of material implication.

David Lewis (1973) and Robert Stalnaker (1968) responded, as Frank Jackson points out in chapter 11, by proposing that (A → B) is true if and only if the closest A-world, the possible world most like the actual world at which A is true, is a B-world, a possible world at which B is also true. This account is attractive because it avoids the paradoxes of material implication, while making Modus Ponens and Modus Tollens valid, as is intuitively correct, and explaining why Strengthening the Antecedent, Transitivity, and Contraposition are invalid. A different response to the paradoxes is the no-truth theory, which states that conditionals have justified assertion or acceptability conditions but not truth conditions.

Vagueness

There is a philosophical *problem* of vagueness because of the sorites paradox, an instance of which is the following inference: (1) A person with $50 million is rich. (2) For any n, if a person with $n is rich, then so is a person with $n − 1¢. (3) Therefore, a person with only 37¢ is rich. This inference constitutes a *paradox* because it appears to be valid, each of its two premises appears to be true (at least when considered on its own), and the conclusion certainly appears to be false. All theorists recognize that the weak link in the inference is the "sorites premise," (2), but they disagree as to what exactly is wrong with that premise and the intuitively compelling argument for it.

In chapter 12, Stephen Schiffer reviews the best known attempts to account for vagueness, and thereby to solve the sorites, and finds them all problematic. He argues that vagueness is neither an epistemic nor a semantic notion, but rather a *psychological* notion, one explicable in terms of a previously unnoticed kind of partial belief he calls vagueness-related partial belief, and which he contrasts with the familiar kind of partial belief, which he calls standard partial belief, that is generally assumed to be normatively governed by the axioms of probability theory. Bringing his psychological account of vagueness to bear on the sorites, Schiffer argues that the paradox doesn't have the sort of neat solution theorists of vagueness typically seek, but instead admits of no determinate resolution.

The semantics of non-factualism, non-cognitivism, and quasi-realism

Non-factualism in some particular area of speech is the claim that the sentences in that area do not function purely representatively, expressing beliefs, but rather express some other mental states that a speaker is voicing. Simon Blackburn lists three motivations for this non-factualism in chapter 13. The metaphysical motivation may be the most important: the apparent queerness of any facts that would be represented. Thus, the apparent queerness of moral facts has led many to non-

factualism, or "non-cognitivism," about moral discourse. Non-factualism has been proposed in many areas including causation and religion.

Peter Geach (1962, 1965) raised a severe problem for non-factualism. How can it account for 'indirect' contexts, ones where a sentence is used, but not asserted or put forward as true. Consider, for example, the role of 'lying is wrong' in the statement 'If lying is wrong, then getting your little brother to lie is wrong.' That role is surely not the expression of an attitude. This problem leads to another: preserving the validity of an argument from that conditional statement together with 'Lying is wrong' to the conclusion 'Getting your little brother to lie is wrong.' Both Blackburn himself (1984) and Alan Gibbard (1990) have proposed solutions that have been the subject of considerable controversy, as Blackburn brings out.

A debate has arisen about whether non-factualism is compatible with a deflationist view of truth (see chapter 20). That view of truth appears to be an example of non-factualism although some think that it takes away the terms in which non-factualism can be formulated. Blackburn's "quasi-realism" which attempts to mirror everything a realist wants to say whilst not having any realist commitment is one way of responding to this apparent conflict.

Part III: Reference

The revolution in the theory of reference

Reference, in ordinary parlance, is aboutness. "What are you referring to?" is more or less equivalent to "What are you talking about?" If the utterance was an ordinary declarative sentence, it's usually a question of what the *subject* of that sentence is, and investigation into reference has understandably focused on terms in subject position.

Many different terms can occur in subject position – chapters 14–18 examine various ones in turn – so an important question is whether or not one theory of reference fits all. The short answer seems to be no. For instance, J. S. Mill (1843) argued that a proper name like 'Dartmouth' does not depend for its reference on any descriptive associations it may have, but he thought general names like 'horse' had their reference determined by an associated description. Frege (1892), as we have already noted, thought all names were descriptive, and indeed that proper names were equivalent to definite descriptions, a view espoused and defended by Russell (1905, 1919). Their views were orthodox until overturned, largely by Kripke (1980), who argued that Mill was right about proper names and wrong about general names. In the last forty-plus years, philosophers of language have continued this lively exchange, and turned the same critical focus upon other referring expressions, including descriptions themselves, all the while attending to developments in – with the hope of reciprocal illumination – logic and linguistics.

Two main questions arise concerning reference: (1) what is the mechanism by means of which reference is secured? (2) what is the meaning of a referring expression? The discussion of these issues invokes three different distinctions that we have already mentioned in discussing chapter 6. The first is metaphysical: a *necessary* true proposition could not have been false, and a *contingent* true proposition might have been false. The next is epistemic: an *a priori* true proposition is knowable independent of experience, and an *a posteriori* true proposition is not. The third is semantic: an *analytic* true proposition is true in virtue of meaning, and a *synthetic* true proposition is not.

Kripkean arguments show to the satisfaction of most that a pure description theory of names is inadequate to answer either the mechanism question or the meaning question. As for mechanism, the view that speakers succeed in referring in virtue of knowing a uniquely identifying – yet non-circular – set of descriptions is beset by ignorance and error problems, as William Lycan notes in chapter 14. As for meaning, archetypal descriptions and archetypal names just seem too different. For example, let 'the F' be the description alleged to constitute the meaning of the name 'Aristotle' for speaker S. Then, according to the description theory, 'Aristotle is F' should be analytic, known a priori and necessary (provided Aristotle exists). Intuitively, however, the sentence seems as synthetic, a posteriori, and contingent as they come.

It's just implausible that competent speakers must have essential properties of an individual in mind in order to refer to it, and similar considerations apply to natural kinds, as Stephen Schwartz notes in chapter 15, drawing on Kripke and Hilary Putnam (1975). For instance, reference to gold and water succeeded long before the a posteriori discovery of the molecular structures that are their essential properties. And the essential properties of biological kinds are even now far from settled.

The works of Kripke and Putnam offer an important competitor to the description theory – one that is capable of explaining the relative ease of referring despite ignorance and error – the causal or historical theory. As to mechanism, the theory distinguishes reference *grounding* from reference *borrowing*. Grounders of a name are relative experts, in more or less direct contact with the thing named, but others borrow the reference successfully when their tokenings are appropriately causally related to those of grounders.

As to meaning, many have combined the causal theory of mechanism with a Millian theory of meaning, called "the direct reference theory," according to which there is no content to a name over and above its reference. There is another option, however; the one mentioned in discussing propositional attitude ascriptions: one might take the meaning of a name to be its particular causal mode of referring to its bearer; the causal network underlying the name determines its meaning (Devitt and Sterelny 1999).

A pure causal theory is anyway inadequate as a theory of mechanism, thanks to the *qua* problem. In the case of natural kind terms, a grounder is in contact with an individual that is a token of several different kinds. In order to secure reference to

just one kind, the grounder's intentions matter, and so the description theory turns out to contain a grain of truth after all. This suggests a third account, a hybrid descriptive/causal theory. Lycan sets out problems for all three accounts and the varieties of associated theories of meaning. Schwartz argues that even if a hybrid account succeeds for natural kind terms, other general terms are less amenable; think particularly of "artifactual" kind terms like 'chair' and social kind terms like 'philosopher.'

Schwartz brings out an important legacy of the "causal revolution" in the theory of reference. This is semantic externalism: a person's relations to the world, including her social world, have a big role in determining the meanings of her terms. A remark of Putnam is the slogan for this externalism: "meanings just ain't in the *head*" (1975: 227). We shall have more to say about this below.

Descriptions

Descriptions come in two basic surface forms: definite and indefinite. Russell's 1905 account of definite descriptions analyses them as object-independent, non-referring terms. In chapter 16, Peter Ludlow and Stephen Neale canvass the many challenges to the Russellian view from referentialists, who argue that at least some definite descriptions – those appearing in what Donnellan (1966) called referential uses – are indeed referring terms. Some referentialists extend the claim to indefinite descriptions as well. Russellians respond by leaning heavily on the distinction between what is literally said, and what is pragmatically communicated, in order to preserve a semantic unity in definites. Ludlow and Neale make the case for Russellianism, and consider the view that there is no semantic distinction between definites and indefinites, either. Whatever the outcome of these debates, Russell's Theory of Descriptions remains an outstanding contribution to philosophy of language. In particular, it has provided a productive framework for philosophers to think about the role of quantification in language.

Indexicals

Pragmatic considerations also weigh heavily in chapter 17, in John Perry's discussion of indexical terms such as "I," "here," and the demonstrative "this." These terms vary their reference according to the context of use, but how do they do this, and why? Perry argues that the meaning of an indexical term is a property of the expression type, which together with the relevant particular context determines the content of an utterance. So "This is Tuesday" has the same meaning whenever it is uttered, but varies in content depending upon what day it is uttered. In considering such variations, Perry makes two distinctions. Is the designation automatic as with 'I' or does it vary according to the intentions of the speaker as with 'now'? Does the reference depend only on the speaker, time, or place of

utterance as with 'now' or 'here,' or does it depend on other facts as with 'this'? Why do we use indexicals? In Perry's view, "to help the audience find supplementary, utterance-independent, channels of information about" the object referred to.

Anaphora

As Perry notes, a pronoun can function like a demonstrative; for example, 'he' in "He loves Sheila," said with a gesture toward Ralph. But pronouns have important other uses brought out in Neale's chapter 18. They can function as bound variables; for example, 'he' in "Every man loves Sheila but he is always disappointed." And they can be anaphoric; for example, 'he', 'she', and 'him' in "Ralph loves Sheila, and he thinks she loves him back." Anaphoric pronouns are the focus of Stephen Neale's discussion in chapter 18.

Neale shows first that pronouns exhibit a variety of behaviors that appear to distinguish them from the bound variables of formal logic. One response to these phenomena is to posit a systematic ambiguity between bound and indexical uses. Another (a methodology paralleling that of the Russellians discussed in chapter 16), is to sweep them up in a pragmatic theory. A third is to regard anaphoric pronouns as standing proxy for descriptions. Neale proposes speaking neutrally of the *binding* of anaphoric pronouns without commitment to one or other of these approaches.

Neale discusses these options in historical context, in which Chomsky's Binding Theory has played a central role. The theory aims to provide syntactic constraints on interpretation, but binding cannot, argues Neale, be purely syntactic. After examining a welter of examples from linguistics, Neale concludes that pronouns do not function as bound variables, but rather *contain* variables that may be bound, a version of the descriptive approach.

Naturalistic theories of reference

The final two chapters concern what might be called metasemantic issues concerning reference, and so dovetail with the concerns of parts I and II. In chapter 19, Karen Neander examines a methodological program for thinking about reference: naturalism. An obvious first step in a naturalistic program is to explain linguistic phenomena such as reference in terms of mental phenomena. Hence naturalistic theories of reference have focused on mental representations.

Causal theories of mental content gain impetus from thought experiments like Putnam's (1975) about Twin-Earth, a planet where all the waterish stimuli consist of XYZ rather than H_2O. Earthling "water" thoughts seem to be about H_2O and not XYZ, even though (a couple of centuries back, at least) nothing strictly in the mind/brain of Earthlings determines that this is so: no images, associated descrip-

tions, or whatever, determine that we refer to H_2O rather than XYZ. This suggests that reference is at least partly determined by causal relations to things external to the thinker; it suggests semantic externalism. These causal relations to particular environmental features make nice naturalistic candidates for reference determination. But a "crude" causal theory faces insurmountable difficulties, many involving a failure to distinguish the "right" causings of mental tokens from the "wrong" ones. Attempts to supplement the causal account to overcome these difficulties include Fodor's asymmetric dependence theory, in which the wrong depend on the right, but not vice versa; teleosemantic theories, which postulates *functions*, understood as what items were selected for, as the arbiter of right and wrong causings; and information theories, according to which representations carry information in virtue of the causal regularities they participate in. Even though none of these approaches seems to solve all the difficulties, they present promising ideas for naturalizing reference.

Truth

Some common notions can appear very mysterious upon inspection. Truth seems obvious and familiar, but what is it? The central issue is whether or not " 'p' is true" says anything more than "p." If so, what kind of property is the property of being true? Will it comport with naturalistic theories of meaning and reference? Vann McGee centers the discussion in Chapter 20 upon a theory that attempts to define truth in non-semantic terms. The theory in question is Tarski's, which enables us to derive, for each sentence s of a language L, a "T-sentence" of the form:

s is a true sentence of L iff _____,

where what fills the _____ is the translation into English of s. The famous paradigm is " 'Snow is white' is true in English iff snow is white." Tarski's theory supplies no general definition of truth, but rather only of truth in a language. Furthermore it applies only to a range of formal languages (despite the paradigm). Within these limits, the theory is "an undoubted triumph," says McGee, but the fact that its methods are inapplicable to natural languages – the Liar paradox and its variations led Tarski to conclude that no consistent account of truth in a natural language is possible – raises the threat "that substantial aspects of human life lie forever beyond the reach of our human understanding."

A range of responses to Tarski's restrictions have been contemplated, most notably one by Kripke (1975), motivated by the idea that the Liar Sentence, and others, are neither true nor false. Aside from this problem, Hartry Field (1972) argued that the theory needs to be supplemented by theories of reference (like those just mentioned) if it is to provide a physicalistically respectable *explanation* of truth; the list-like definitions of reference on which Tarski's theory rests are

insufficient. The theory as it stands does not capture the idea that truth consists in a robust correspondence to the facts. McGee concludes with a discussion of an alternative to the correspondence conception. This is the "deflationary" view that the truth term is a logical device for disquotation; the term does not refer to a robust property that plays an explanatory role in science. The idea goes back at least to Frege (1892), who wrote,

> One can, indeed, say: 'The thought, that 5 is a prime number, is true.' But closer examination shows that nothing more has been said than in the simple sentence '5 is a prime number.'

On this view, the addition of the words "is true" to a sentence adds no further content. One battleground between correspondence and deflationary theories is over vagueness (cf. chapter 12).

As we remarked at the outset, hardly anything in the philosophy of language is settled. There is plenty of work to be done, but that does not mean that no progress has been made. Modern analytic philosophy of language has made a substantial contribution to our overall understanding of the world we occupy, and the language we use to talk about it.

Note

Our thanks to Panu Raatikainen for helpful comments on a draft of this introduction

Part I

Foundational Issues