
International Library of Philosophy
Edited by José Luis Bermúdez, Tim Crane and Peter Sullivan

Advisory Board: Jonathan Barnes, Fred Dretske, Frances Kamm,
Brian Leiter, Huw Price and Sydney Shoemaker

Recent titles in the ILP:

The Facts of Causation
D.H. Mellor

The Conceptual Roots of Mathematics
J.R. Lucas

Stream of Consciousness
Barry Dainton

Knowledge and Reference in Empirical Science
Jody Azzouni

Reason Without Freedom
David Owens

The Price of Doubt
N.M.L. Nathan

Matters of Mind
Scott Sturgeon

Logic, Form and Grammar
Peter Long

The Metaphysicians of Meaning
Gideon Makin

Logical Investigations, Vols I & II
Edmund Husserl

Truth Without Objectivity
Max Kölbel

Departing from Frege
Mark Sainsbury

Truth Without Objectivity

Max Kölbel



London and New York

International Library of Philosophy
Edited by José Luis Bermúdez, Tim Crane and Peter Sullivan

Advisory Board: Jonathan Barnes, Fred Dretske, Frances Kamm,
Brian Leiter, Huw Price and Sydney Shoemaker

Recent titles in the ILP:

The Facts of Causation
D.H. Mellor

The Conceptual Roots of Mathematics
J.R. Lucas

Stream of Consciousness
Barry Dainton

Knowledge and Reference in Empirical Science
Jody Azzouni

Reason Without Freedom
David Owens

The Price of Doubt
N.M.L. Nathan

Matters of Mind
Scott Sturgeon

Logic, Form and Grammar
Peter Long

The Metaphysicians of Meaning
Gideon Makin


Logical Investigations, Vols I & II
Edmund Husserl

Truth Without Objectivity
Max Kölbel

Departing from Frege
Mark Sainsbury

Truth Without Objectivity

Max Kölbel



London and New York

Truth Without Objectivity

'This is the most interesting, carefully constructed, and challenging exposition and defence of the – initially repugnant – view that truth is relative which I have read. Whilst written with admirable simplicity and clarity, its argumentation has both philosophical depth and great subtlety.'

Bob Hale, University of Glasgow

According to the mainstream view in the philosophy of language, to know the meaning of a sentence is to know the conditions under which the sentence would be true. The view, however, is challenged by non-objective sentences such as sentences on matters of taste or value: these do not appear to be either true or false, but are generally taken to be meaningful. How can this conflict be resolved?

Truth Without Objectivity examines different ways of resolving this fundamental problem, before developing and defending a relativist theory of truth. Standard solutions maintain either that in uttering non-objective sentences speakers make implicit reference to their own preferences and thus have unproblematic truth conditions, or that they have no truth conditions at all. Max Kölbel argues that both of these proposed solutions are inadequate, and that the third well-known position, minimalism, can only solve the problem if it is developed in the direction of relativism about truth.

Kölbel defends the idea that truth (as invoked in semantics) is a neutral notion: a sentence's possessing a truth condition does not yet entail that it concerns an objective subject matter, because truth and objectivity are independent of one another. He argues that this notion of 'truth without objectivity' leads directly to a relativist theory of truth, and goes on to defend his form of relativism against the usual objections to such theories.

Truth Without Objectivity is a valuable and clear discussion of one of the most interesting questions in philosophy. It will be of particular interest to all students of the philosophy of language, metaphysics and metaethics.

Max Kölbel is Lecturer in Philosophy at the University of Birmingham.

First published 2002
by Routledge
11 New Fetter Lane, London, EC4P 4EE

Simultaneously published in the USA and Canada
by Routledge
29 West 35th Street, New York, NY 10001

Routledge is an imprint of the Taylor & Francis Group

© 2002 Max Kölbel

The right of Max Kölbel to be identified as the author of this work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1999.

Typeset in Times by Max Kölbel
Printed and bound in Malta by Gutenberg Press Ltd

All rights reserved. No part of this book may be reprinted or reproduced or utilised in any form or by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying and recording, or in any information storage or retrieval system, without permission in writing from the publishers.

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication Data
has been applied for

ISBN 0-415-27244-0 (hbk)
ISBN 0-415-27245-9 (pbk)

Contents

<i>Preface</i>	ix
<i>Introduction</i>	xii
1 Truth-Conditional Semantics	1
1 One version of truth-conditional semantics	2
2 Compositionality	3
3 Force indicators and content indicators	4
4 Truth definitions 'serving as' meaning theories	6
5 Application to natural languages	10
6 Proper communicative function	10
7 Communicative acts	14
8 Summary and conclusion	16
2 Excess Objectivity	18
1 Excess objectivity	18
2 Three ways of solving the problem of excess objectivity	20
3 Kinds of objectivity	21
4 Wright's criterion of cognitive command	22
5 The problem of a priori error	24
6 Does plain error amount to cognitive failure?	25
7 A new criterion for objectivity	28
8 How to solve the problem of a priori error	31

3 Revisionism	35
1 How to revise a content ascription	35
2 The attraction of revisionism	36
3 Objection to revisionism: denial	38
4 Revising revisionism	40
4 Expressivism	43
1 Some expressivists	44
2 Unendorsed contexts	46
<i>Searle's adequacy condition</i>	47
<i>Geach and his Frege-point</i>	49
3 Two ways of meeting Geach's challenge	52
<i>The force-indicator approach</i>	53
<i>The content-indicator approach</i>	57
4 Hale and the grammatical uniformity of declarative sentences	60
5 Resolving Hale's problem	64
6 Concluding remarks	67
5 Soft Truth	69
1 The traditional debate	70
2 Two notions of truth	72
3 Davidson's programme	74
<i>The biconditional doctrine</i>	76
<i>How to derive intensional theorems</i>	78
<i>Davidson's programme before the truth doctrine</i>	81
<i>The emergence of the truth doctrine</i>	82
4 Radical interpretation without <i>the</i> notion of truth	85
5 'Meaning as truth condition'	88

6 Relative Truth and Linguistic Communication	91
1 An absolutist model of linguistic communication	92
<i>Background assumptions</i>	92
<i>Simple cases of communication</i>	93
<i>Cases of incredulity</i>	94
2 Disagreement without error	98
3 The relativized truth framework	100
4 The explanatory status of the theory	103
5 Methods of belief acquisition	105
6 Extension of the framework: shades of objectivity	107
7 One further application: fiction	109
8 Alan Gibbard: revisionist, expressivist or relativist?	110
7 Defence of Relativism	116
1 The essence of relativism	116
2 Classifying relativisms	118
3 Newton-Smith and the traditional conception of proposition	119
4 The dilemma/self-refutation	122
<i>Horn (ii) contradicts (GR)</i>	122
<i>Horn (ii) is pragmatically self-refuting</i>	123
<i>Given horn (ii), it is pointless to assert (GR)</i>	123
5 The possibility of communication	124
6 Problems for Protagorean relativism	126
<i>Notes</i>	129
<i>Bibliography</i>	142
<i>Index</i>	147

Preface

This book has taken a long time to mature. I first started thinking about the metaphysical consequences of truth-conditional semantics in 1993, when I was a postgraduate student at King's College London (KCL). Originally I thought truth-conditional semantics was misguided because of these consequences, and in my MPhil thesis I therefore defended a form of expressivism that leads away from truth-conditional semantics. Later on, however, I realized that truth-conditional semantics is a kind of Kuhnian paradigm. Most theorists who work in semantics (in philosophy, artificial intelligence, computer science, linguistics, cognitive science, psychology) work within this framework. In fact, for many theorists 'semantics' just means truth-conditional semantics. So I started exploring the advantages of truth-conditional semantics, and enquiring whether these advantages couldn't be had without the unwanted metaphysical consequences.

This book is the outcome of that enquiry. An intermediate step was my PhD thesis, which I submitted in September 1996, and which defended the same thesis as this book, under the same title. Because of the kind encouragement of Peter Sullivan and José Bermúdez, I started rethinking and rearranging my material with a book in mind. For various reasons, however, it took five years to complete. First, there were many difficult questions about which I felt I needed more time to think. I no doubt felt this way partly because of the formidable objections made by colleagues whenever I presented my work. Another reason for it taking so long was my (mostly involuntary) nomadic lifestyle: during those five years I lived and worked in five different places and moved house even more often. It is hard to say what overall effect the instability of location had on my frame of mind and on the book. One effect, however, was without doubt desirable: my work got exposed to the criticism of a wider variety of people than it would otherwise have been.

For the same reason, the list of people I would like to thank for their input to this book is very long. It includes my exceptionally generous PhD supervisors, teachers, fellow students and friends from my days as a KCL graduate student. It also includes colleagues and friends from the fabulous Instituto de Investigaciones Filosóficas at the Universidad Nacional Autónoma de México (UNAM)

in Mexico City, where I spent two unforgettable years. It further includes colleagues from Christopher Habel's Graduiertenkolleg Kognitionswissenschaft at Hamburg University, where I spent most of 1999. Then there are those of my colleagues, students and friends at Swansea University, who made it possible for me to survive a turbulent year there, and to remember occasionally that I was meant to be working on this book. The list also includes students, colleagues and friends from Cambridge, where I spent the last year in comparative peace and in a stimulating intellectual atmosphere. There are also the many colleagues who have commented on my work in conferences, workshops, talks and the like, especially the two Hamburg-Barcelona Workshops in October 1999 and February 2000, and the Workshop on Belief and Meaning at Regensburg in May 2000. Finally, the list includes my long-standing friends, companions and family, in particular my brother and his family, who allowed me to prepare the penultimate version in their happy company in the summer of 2000.

I warmly thank all the people described above for whatever help they extended to me. What follows is an inadequate attempt to compile the names of most of them:

Arif Ahmed, Brom Anderson, Helen Baldwin, Stephen Barker, José Bermúdez, Simon Blackburn, Johannes Brandl, the British Academy, Krister Bykvist, Marta Campdelacreu, Mat Carmody, Donald Davidson, Peter Dunstan, Rebecca Dunstan-Harvey, Dorothy Edgington, Anne Edwards, Maite Ezcurdia, Miguel-Angel Fernández, Michalis Filippou, Lorena García, Manuel García-Carpintero, Laurence Goldstein, Mario Gómez-Torrente, Aruna Handa, Wolfram Hinzen, Shirin Homann, James Hopkins, Keith Hossack, Guillermo Hurtado, Carrie Jenkins, Peter Kail, Hans Kamp, Ariel Kernberg, Jonathan Knowles, Friederike Kölbel, Johannes Kölbel, Reinhard Kölbel, Sebastian Kölbel, Tilo Kölbel, Martin Kretschmer, Sharon Kretschmer, Wolfgang Künne, Hallvard Lillehammer, Ieuan Lloyd, Dan López de Sa, Fraser McBride, M.M. McCabe, Greg McCulloch, John McDowell, Josep Maciá, Hugh Mellor, the Mind Association, Raymundo Morado, Howard Mounce, the Fellows of New Hall, Cambridge, Samir Okasha, Alex Oliver, Eric Olson, Cathy Osborne, Karla Otero, Martin Otto, David Papineau, Michael Potter, José Antonio Robles, Alexandro Rojas, Rikardo Rojas, Hans Rott, Mark Sainsbury, Ricardo Salles, Katja Sengelmann, Linda Sengelmann, Ruben Sengelmann, Mark Siebel, Wolfgang Spohn, Christian Stein, Pedro Stepanenko, Peter Sullivan, Rupert Summerton, Mark Textor, Alan Thomas, Ian Tipton, Mario von der Ruhr, Neal Wagstaff, Denis Walsh, Ginny Watkins, Ruth Weintraub, Bernhard Weiss and Ieuan Williams.

I would like to thank the editor of *Theoria* for permission to use, in Chapter 2, material originally published as 'A Criterion of Objectivity', in *Theoria: Revista de Teoría, Historia y Fundamentos de la Ciencia* 15 (May 2000), pp. 209–28.

I would like to thank the editors of *Crítica: Revista Hispanoamericana de Filosofía* for permission to use, in Chapter 4, material first published as

'Expressivism and the Syntactic Uniformity of Declarative Sentences' in *Crítica* 29, no. 87 (December 1997), pp. 3–51.

Finally, I would like to thank the editor of *The Journal of Philosophy* for permission to use, in Chapter 5, material originally published as 'Two Dogmas of Davidsonian Semantics' in *The Journal of Philosophy* 97 (December 2001), pp. 613–35.

MK

August 2001

Introduction

There is a certain approach to theorizing about language that is called 'truth-conditional semantics'. The underlying idea of truth-conditional semantics is often summarized as the idea that the meaning of a sentence can be specified by giving the condition under which it would be true. This is then condensed into the slogan that the meaning of a sentence is its truth condition. The slogan has intuitive appeal, because the meaning of a sentence is what one knows when one understands it, and it is plausible to suppose that knowing the conditions under which a sentence is true is to understand it.¹

If this characterization is correct, then truth-conditional semantics faces a huge problem. It seems to presuppose that all sentences are evaluable as true or false, but there are many reasons to think that not all sentences are candidates for truth or falsehood. This book is about one kind of reason for doubting that all sentences have truth conditions: non-objectivity.² Many philosophers believe that, for example, values and probabilities aren't objective. On this view, if a (declarative) sentence concerns a matter of taste, a moral question or the probability of an event, then it concerns something non-objective. But how can such a sentence then be evaluated as true or false? It seems that truth-conditional semantics rests on a highly dubious presupposition.

This book's aim is to examine possible solutions to this problem. Should truth-conditional semanticists insist that all sentences nevertheless have truth conditions, and if so, does this entail that all sentences describe objective reality? Or should non-objective sentences be exempted from the truth-conditional treatment? These questions have received some attention from metaphysicians and meta-ethicists and they continue to be debated. By contrast, they have been largely neglected by natural language semanticists who work within the truth-conditional paradigm, i.e. by those who attempt to make the truth-conditional approach work for particular natural language constructions. This is surprising. According to the truth-conditional slogan, the meaning of a sentence is its truth condition. So isn't the question whether, say, evaluative sentences are truth evaluable an obvious and urgent preliminary question each truth-conditional semanticist has to settle?

The neglect can be explained. Behind the truth-conditional slogan is a very complex view which has little, if anything, to do with the profound metaphysical significance the slogan suggests. Truth-conditional semantics is a view concerning the *form* a theory of meaning (semantic theory) for a particular natural language should take, namely the view that it should take the form of an axiomatic theory whose recursive axioms generate theorems of the form '*s* is true iff *p*' for every sentence *s* of the language in question. This view has guided the work of many philosophers of language and linguists. But its primary motivation is not the idea that a theory of meaning for a language ought to tell us something about how that language's expressions relate to extra-linguistic reality. Rather, it is motivated by the need to describe in a precise way, how the meanings of complex expressions depend on the meanings of their parts. Most theorists who work within the truth-conditional paradigm do so because it allows them to account for the compositionality of languages, be this because they want to explain learnability, because they are interested in the logical properties of a language or for some other reason. The attraction of truth-conditional semantics lies not in the fact that it connects the notion of meaning with the notion of truth, but rather that in so doing it can map out the semantic structure of particular languages.

My stance on the issue of the truth evaluability of non-objective sentences is driven by the view of truth-conditional semantics I just outlined. Semantic theory is primarily concerned with the phenomena of language, not with metaphysics. Thus, if certain metaphysical intuitions about the objectivity of truth threaten the otherwise healthy project of constructing semantic theories that employ a notion of truth, then the semanticist should, if possible, deny that the notion employed in semantics is identical with the notion whose metaphysical features create the problems. This book defends the view that the notion of truth employed in semantic theories is a metaphysically neutral notion, according to which a sentence's possessing a truth condition does not yet entail that it concerns an objective subject matter. Truth in semantics is truth without objectivity.

Nevertheless, it would be misleading to describe this book as metaphysics-free. I argue that the notion of truth invoked in semantics can be identified with the notion of truth we actually employ—or rather with one of the two notions we employ. I also argue that a notion of truth without objectivity must be a notion of relative truth. Finally, I defend relativism about truth (and other forms of relativism) against the charge of incoherence. Thus, while the semanticist should be allowed to operate with his or her own semantic notion of truth without interference from metaphysics, I believe that there is a coherent metaphysical theory of truth that serves the purposes of semantics. Readers who are not prepared to agree with me on these metaphysical points can still agree with me in the philosophy of language concerning the notion of truth used in semantics. The metaphysics in this book can be separated from the semantics.

My argument for these conclusions is not as quick as the previous two paragraphs suggest. I start by explaining in Chapter 1 what truth-conditional semantics is, or rather how I view it. In Chapter 2, I set up the problem: truth-

conditional semantics seems to presuppose that all sentences are truth evaluable and this conflicts with the common view that some sentences aren't truth evaluable because they concern non-objective matters. I distinguish three ways of dealing with this problem: revising the contents assigned to sentences (revisionism), exempting some sentences from truth-conditional treatment (expressivism), and employing a notion of truth that does not entail objectivity (soft truth). I then devise a criterion for objectivity which is inspired by, but in crucial respects different from, Crispin Wright's criterion of 'cognitive command'. According to this criterion, a proposition is objective if a mere disagreement on that proposition shows that a mistake has been made. I argue that the only way in which a proposition could be truth evaluable yet non-objective is that truth is relative. Relativism about truth is therefore a consequence of the soft-truth strategy.

Radical solutions are only palatable if no less radical alternative is available. That's why in Chapters 3 and 4 I examine revisionism and expressivism in some detail. Revisionism is the claim that sentences on non-objective matters are generally elliptical, and involve an implicit indexical element. For example, 'laver bread is tasty' might be claimed to be elliptical for 'I find laver bread tasty'. I argue that any such claim is false, because there are demonstrable differences in meaning between the original sentences and the ones that they are said to be elliptical for.

Expressivism requires a much more detailed examination. Expressivists want to exempt problematic sentences from truth-conditional semantic treatment and account for the meaning of these sentences in some other way, usually claiming that they exhibit some special kind of illocutionary force. They face the problem that the objectivity or non-objectivity of a sentence's subject matter has usually little influence on that sentence's syntactic properties. That is, the sentences the expressivist wants to exempt from truth-conditional treatment can be combined with, and embedded in, other sentences—even those that have been approved for the truth-conditional treatment. In short, the expressivist is up against the syntactic uniformity of sentences that aren't uniform in respect of their objectivity status. I argue that even though sophisticated expressivists might overcome these difficulties, they will end up with a uniform *non*-truth-conditional semantics. Thus, even though expressivism can be a coherent position, it is then no longer a solution to the problem I posed, i.e. not a solution for a fundamental problem *within* truth-conditional semantics.

With revisionism and expressivism discarded as solutions to the problem, I move on to defend my own view in Chapters 5, 6 and 7. First, I show in Chapter 5 how truth-conditional semantics took a wrong turn in the early 1970s, when Davidson started claiming that the notion of truth plays a crucial explanatory role in semantics. I argue, inspired by McDowell's writings from the 1970s, that within the semantics of natural language, truth should be regarded as a theoretical notion that can be fully understood by its role in a semantic theory. I also

argue that another Davidsonian dogma, the view that a theory of meaning can only generate extensional theorems, is unjustified.

The path is then clear for my positive account of relative truth in chapter 6. I expound a theory according to which (1) truth is relative to *perspectives*, (2) each thinker *possesses* a perspective and (3) a thinker ought not to believe anything that isn't true in his or her own perspective. A perspective, on this theory, is just a function that evaluates all propositions consistently; thus there is nothing philosophically substantial in the postulation of perspectives. The substantial element of the theory is rather the claim that the relation of perspective possession is constrained in a certain way by certain *a priori* norms of communication. These norms specify that in some areas of discourse disagreements indicate that a mistake has been made and that therefore reasoned discussion would be worth while. These are areas one might call objective. Thus objectivity, on this view, is the result of certain rules of communication. I also show how this theory of perspective possession can be further refined to make finer distinctions between more or less objective topic areas.

The final chapter examines relativism in general. This is necessary because discussions of relativism are often hampered by a lack of clarity concerning the nature of relativism. Often one form of relativism is dismissed because of problems that arise only for other forms of relativism. I therefore begin the chapter by developing a scheme of classification for different forms of relativism. After that, I use the classificatory scheme to examine how well-known objections to relativism fare against the various forms of relativism. The result of this examination is that the impact of these objections is surprisingly small.

The main conclusions I shall reach thus are: (1) the best response to the problem of non-objective sentences in truth-conditional semantics is the adoption of a truth notion on which mere truth evaluability does not yet entail objectivity; (2) any truth notion to fit that bill must be relative; (3) an independently fruitful theory of perspectives and perspective possession can be devised; and (4) this form of relativism, even though it is a global form, does not fall to any of the usual objections.

Truth-Conditional Semantics

Humans use language to communicate. They utter certain sounds and respond to others uttering sounds in ways that constitute communication. What exactly is going on? How is this possible? We all know that it is because the expressions of a language are *meaningful* that we can use them for communication. Philosophers of language try to expand on this tautological and uninformative answer by constructing a more detailed theory of meaning. Such a theory may answer questions such as 'What is it for an utterance to be meaningful?', 'Why does a particular sound have the meaning it has?', 'How do people know what an utterance means?' and so on. By answering such questions, and by introducing new concepts and principles governing these concepts as well as linking them with other, familiar concepts, a theory of meaning will provide a more detailed explanation of what is going on when humans use language to communicate.

Evidently this description of the task of the philosophy of language leaves room for a very wide variety of theoretical activity, according to the different aspects of language use a theorist might be interested in. Nevertheless, one broad style of theorizing about linguistic meaning currently dominates the philosophy of language (and other disciplines that theorize about linguistic meaning, such as theoretical linguistics and artificial intelligence). I would like to call this common approach 'truth-conditional semantics'. Theorists working within the truth-conditional paradigm are united by the role their theories assign to a notion of truth or truth condition. Very roughly, these theorists believe that it is advantageous to characterize the meaning of a sentence (or at least the central element of its meaning) by giving the condition under which the sentence would be true. The meaning of a subsentential expression then consists just in the contribution it makes to the meaning of sentences in which it occurs. As I shall explain in more detail below, one great advantage and original motivation for this approach is that it provides a way of describing systematically how the meaning of complex expressions depends on the meanings of their simpler constituents.

Any theory of meaning within the truth-conditional paradigm faces a problem: it makes the contentious assumption that every meaningful sentence (or its content) is a candidate for truth or falsehood. There is, of course, an obvious reason why not every sentence is truth-evaluable: some sentences are there to enable us

to ask questions or to issue commands, such as 'Is he in?' or 'Leave me alone!'. One would not normally describe such sentences (or utterances of them) as true or false. This is not the problem I have in mind, nor is it an irresolvable problem for a truth-conditional theorist. Typically the theorist will restrict his or her thesis that sentences have truth conditions to declarative sentences only and characterize the meaning of non-declaratives indirectly, e.g. by recourse to corresponding declaratives. Alternatively, the theorist can reformulate the thesis as one about sentence *contents*, so that declaratives and interrogatives alike have truth-evaluable contents. Declarative and interrogative sentences (as opposed to their contents) are then viewed as truth-evaluable only in a derivative sense, if at all. I shall discuss one way of doing this shortly.

But once interrogatives, imperatives etc. are taken care of there is another, more difficult problem with the assumption that all sentences (or their contents) are truth-evaluable. I have in mind declarative sentences that concern matters that are not objective, or for which it is at least debatable whether they are objective. One example are sentences about matters of taste such as 'Johnny Depp is more handsome than Brad Pitt.'¹ Assuming that such a sentence, or what it says, can be true or false arguably amounts to saying that it is an objective matter whether Depp is more handsome or not. Can the truth-conditional theorist avoid this metaphysical consequence of his or her theory?

This book is an extended attempt to answer the question of how these and similarly problematic sentences should be treated in a truth-conditional theory of meaning. Competently answering it requires a theoretically more sophisticated formulation of the problem. Thus I shall begin my investigation by expounding a truth-conditional account of meaning in this chapter, and continue in the next chapter by setting up the problem in a theoretically tractable manner.

1 One version of truth-conditional semantics

Accounts of meaning within the truth-conditional paradigm vary considerably. Some differences reflect substantial disagreements, others reflect merely a difference in emphasis—after all, different theories often address different problems. Yet other differences are merely terminological. Consequently, the account I shall expound is just one truth-conditional account among many, and not all truth-conditional theorists will agree with every detail. However, one particular elaboration of a truth-conditional theory of meaning is needed as a concrete theoretical and terminological framework within which to address the problem of sentences about non-objective matters. My elaboration will be sufficiently broadminded to allow most truth-conditionalists to follow the discussion without excessive interpretative charity. More importantly, I believe that my considerations on the problem of non-objective contents apply to my own account of meaning as much as to any other account within the truth-conditional paradigm.

My account will integrate a truth-conditional theory of content within an overall theory of use. A theory of linguistic meaning ought to explain how it is pos-

sible to use expressions of a language to communicate. In order to arrive at such a very general explanation, it is useful to consider what would be required to describe the meaning properties of all the expressions of just one particular language, i.e. to construct a theory of meaning for a particular language. Such a theory would explain how it is possible for the users of the language in question to use their language for communication, because such a theory would provide a formal model of the knowledge of language users that enables them to use language (whether or not that model is psychologically accurate). But, more importantly, if we knew how in principle such a theory of meaning for a particular language can be constructed, then we would be in a position to explain generally how it is possible to communicate through language (whichever particular language). Later on in this chapter, I shall formulate a theory of meaning for one particular simple language, and this theory shall serve as a simplified model for theories of meaning for natural languages in general.

2 Compositionality

One feature of language is a key motivation for the truth-conditional paradigm: compositionality. Natural languages are compositional in the sense that they allow the composition of compound expressions out of simpler ones in such a way that the meaning of the compound is determined by the meanings of the simpler expressions and their manner of composition. Compositionality seems to be one of the most obvious features of language and can be illustrated by countless examples. What 'wooden box' means obviously depends on what 'wooden' means, on what 'box' means and on how the two words are combined. If 'box' didn't mean what it does, 'wooden box' wouldn't mean what it does. To deny this would be crazy. Could the meaning of any complex expression depend on anything apart from its components and the way they are compounded? Strictly speaking, there could be syntactically complex expressions that are nevertheless not semantically complex. There could be a syntax that provides numerous ways of constructing new expressions without it yet being settled what these new expressions are to mean. Then the meaning of the (syntactic) components (if they have any) would not determine the meaning of the (syntactic) compound.² All this shows is that there may in some cases be syntactic complexity that is semantically inert. But this cannot hide the fact that in countless cases the meaning of compounds is compositionally determined.

Considerations about language understanding and language learning also confirm that compositionality is a pervasive feature of natural languages. Language users can understand and construct expressions they have never used or even heard before. For example, the vast majority of the readers of this book will never before have encountered the following sentence: 'My niece collects snails in a wooden box.'³ Nevertheless, all moderately competent users of English (including, presumably, all my readers) know the meaning of this sentence. It doesn't take much ingenuity to explain this. The sentence is constructed from

words, and in ways, with which competent speakers are familiar. The meaning of the sentence is determined by the meaning of its constituent words and the way in which these constituents are put together. Competent users know (at least implicitly) what the words mean and how they influence the meaning of the whole sentence, given that they are put together in that particular way.³

This observation puts a constraint on theories of meaning for natural languages. Any such theory should explain how exactly the meaning of complex expressions is determined by the meanings of their simpler constituents and the manner of composition. In one word, a theory of meaning for such a language should be *compositional* in order to do justice to the fact of linguistic productivity.

Given the compositionality constraint, we know that a theory of meaning for a natural language cannot just list all the expressions of the language and specify their meaning one by one. Rather, it must first systematically outline all the ways in which expressions can be compounded from a number of primitive elements in a number of admissible ways (syntax). Then it must show how the meaning of compound expressions depends on their compositional make-up by specifying for each mode of composition, and for each primitive element, how they determine the meaning of the expression compounded. As a result, the theory should allow one to derive for every expression of the language what it means. In other words, the theory should have axiomatic structure.

3 Force indicators and content indicators

But what sort of information is it that we should expect to be derivable about each expression? How does one state what an expression means? At this point it is useful to introduce the notion of a *communicative act* (or traditionally 'illocutionary act'). By uttering certain expressions or, more accurately, by uttering sentences, language users perform linguistic acts, such as acts of assertion and question. I can assert that Sam smokes by uttering the sentence 'Sam smokes.', or I can ask whether Sam smokes by uttering the sentence 'Does Sam smoke?'. This is not a coincidence. Rather, it is the proper function of these sentences to allow one to do just this. One way, then, of specifying what an expression means is to state what its *proper communicative function* is. For example, one way of saying what the sentence 'Sam smokes.' means is to say that it can be properly used to assert that Sam smokes. The theory of meaning I am going to outline for a sample language will aim to generate meaning specifications of roughly this form. (I shall say more about communicative acts, my notion of proper use, and about the possibility of improper uses in §6 and §7.)

Not every expression, not even every complex expression, of a natural language can be properly used to perform a communicative act. The expression 'does not smoke', for instance, is not a complete sentence, and therefore cannot be properly used to perform any communicative act.⁴ I shall call expressions that can be properly used to perform a communicative act 'sentences'.

The communicative function of a given sentence will depend on recognizable features of the sentence: its semantically significant features. For example, 'Sam smokes.' can be properly used to assert that Sam smokes, because it is made up of certain words in a certain order, and features a certain kind of punctuation (in spoken language: inflection). Among the semantically significant syntactic features of a natural language sentence are therefore features such as (a) its constituent words, (b) its word order, (c) its punctuation, and for spoken language (d) its inflection (i.e. changes in intonation).

A theory of meaning for a language needs to correlate these syntactic features of a sentence with the sentence's semantic features, i.e. its meaning. To do so, it is useful to distinguish between two different elements in the meaning of a sentence. One element is the sentence's *illocutionary force*: a sentence can, for example, be assertoric or interrogative. If it is assertoric, then it is properly usable for asserting something; if it is interrogative, then it is properly usable for asking a question. The other element is the sentence's *content*: for example the sentence 'Sam smokes.' has the content that Sam smokes. Two sentences can have the same content and different force, as for example the two sentences 'The door is shut.' and 'Is the door shut?'. The first can be properly used to assert that the door is shut, while the second can be properly used to ask whether the door is shut. Similarly, two sentences can have the same force and different contents.

The task of a theory of meaning (of specifying the communicative function of each sentence) can be divided into two subtasks. The first subtask is to identify those syntactic features that indicate a sentence's illocutionary force (call these its 'force indicators') and to formulate a rule that states how these force indicators determine a sentence's force. The second subtask is to identify the syntactic features that indicate a sentence's content (call these its 'content indicators') and to formulate rules that state how a sentence's content depends on these content indicators.

The first subtask is comparatively easy. Each sentence has just one illocutionary force, so all one needs to do is find the force indicators of the language in question and define a function that assigns forces to sentences depending on the force indicator they feature. In English, for instance, assertoric force is indicated by a full stop (or the corresponding inflection) and a certain word order. It is the task of applied linguistics to discover the exact way in which assertoricity is indicated. In my sample language, there are only two force-indicating prefixes, '†' (assertoric force) and '?' (interrogative force). Natural languages allow more than these two illocutionary forces, but I shall not here enter the discussion on which these are.

The subtask of assigning forces to sentences is comparatively easy, because force indicators do not, arguably, have semantic structure. Each complete sentence is marked to have one illocutionary force.⁵ Content indication, however, is more complicated. We can construct sentences with highly complex contents. The job of formulating axioms that formulate how the contents of sentences depend on their content indicators is a real problem. Not only are there many

different ways of combining words and sentences (at least in any natural language), but these modes of construction can also be iterated. This makes it necessary to operate with recursive axioms.

4 Truth definitions 'serving as' meaning theories

This is the point where truth-conditional semanticists turn to mathematical logic, more specifically to the work of Alfred Tarski, the founder of model theory. Tarski devised a way of defining a notion of truth for a formal first-order language (Tarski 1956). A Tarskian definition of truth for a language L allows one to prove a so-called 'T-sentence', i.e. a sentence of the form ' s is true if and only if p ', for each sentence s of L so that ' p ' is the translation of s into the language of the definition (the 'metalanguage').

Truth-conditional semanticists hope to exploit Tarski's technique for their own purposes. First, even though a T-sentence does not say what the content of the sentence in question is, it nevertheless allows one to infer what the sentence's content is, at least if it has been derived in a certain way from appropriate semantic axioms. Second, even though Tarski was defining truth for a formal language only, the truth-conditional semanticist hopes to adopt Tarski's technique for natural languages as well.⁶ To get a concrete idea of how this might be done, let's consider a very simple propositional language, $L1$, which features both content indicators and force indicators:

Syntax for $L1$

Vocabulary:

Names:	a, b, c
1-place predicates:	S, D
2-place predicate:	L
Connectives:	$\&, \neg$
Force indicators:	$\vdash, ?$
Brackets:	$(,)$

Definition of sentential phrase:

For all $\Pi, \alpha, \vdash, \beta, \theta, \rho$:

- (1) If Π is a 1-place predicate and α is a name, then ' $\Pi\alpha$ ' is a sentential phrase.⁷
- (2) If Π is a 2-place predicate and α and β are names, then ' $\alpha\Pi\beta$ ' is a sentential phrase.
- (3) If θ and ρ are sentential phrases, then so are ' $\neg\theta$ ' and ' $(\theta \& \rho)$ '.
- (4) Nothing else is a sentential phrase.

Definition of sentence:

For all θ :

- (5) If θ is a sentential phrase, then ' $\vdash\theta$ ' and ' $?\theta$ ' are sentences and nothing else is.

Thus, ' $\neg(Da \& \neg Sa)$ ' would for example be a sentential phrase of $L1$, ' $\vdash\neg(Da \& \neg Sa)$ ' a sentence. The definition of 'sentential phrase' is recursive: it allows repeated application of the two connectives to form a new sentential phrase each time. So there are infinitely many sentential phrases in $L1$. Nevertheless, we have tight control over the sentential phrases of $L1$, for they can be generated only in ways specified by the recursive definition of 'sentential phrase'. It is now the task of a theory of truth for $L1$ to make sure that truth is defined for each sentential phrase. This again is done recursively:

Theory of truth for $L1$

Definition of reference:

- (6) The referent of a is Alfred, the referent of b is Bernard and the referent of c is Carl.

Definition of true:

For all $\alpha, \beta, \theta, \rho$:

- (7) If α is a name, then ' $S\alpha$ ' is true iff the referent of α smokes.
- (8) If α is a name, then ' $D\alpha$ ' is true iff the referent of α drinks.
- (9) If α and β are names, then ' $\alpha L\beta$ ' is true iff the referent of α loves the referent of β .
- (10) If θ is a sentential phrase, then ' $\neg\theta$ ' is true iff it is not the case that θ is true.
- (11) If θ and ρ are sentential phrases, then ' $(\theta \& \rho)$ ' is true iff θ is true and ρ is true.

Given this theory of truth for $L1$, one can prove a T-sentence for each sentential phrase of $L1$. For example, one can prove the T-sentence

- (T1) ' $\neg(Da \& \neg Sa)$ ' is true iff it is not the case that (Alfred drinks and it is not the case that Alfred smokes).

for the previous example sentence ' $\neg(Da \& \neg Sa)$ '. In doing so, one retraces the syntactic operations by which the sentence has been compounded by applying the corresponding semantic axioms.

How can this help in the task of formulating a theory of content, which can then feed into a theory of meaning that specifies the proper communicative

function of each sentence? All the T-sentences do is give a necessary and sufficient condition for the truth of a sentential phrase. The truth-conditional semanticist's idea is that the *right* T-sentences allow one to *read off* what the content of the sentential phrase is. But not every T-sentence allows one to read off the content. To clarify this problem, let's define 'T-sentence' properly.

A *T-sentence* is a metalanguage sentence of the form:

' θ is true iff p '

where ' θ ' is replaced by a description of an object-language sentential phrase, and ' p ' is replaced by a sentence that does not contain any reference to object-language expressions.

Now suppose we wanted to exploit the T-sentences generated by our theory of truth in order to assign communicative functions to sentences by adding one more axiom and an inference rule:

(12) If θ is a sentential phrase, then ' $\vdash \theta$ ' can be properly used to assert the content of θ .

(TM) If ' θ is true iff p ' is a T-sentence, then θ 's content is the content that p .

Then we could derive the following meaning-specifying theorem for ' $\vdash \neg(\text{Da} \ \& \ \neg\text{Sa})$ ':

(M1) ' $\vdash \neg(\text{Da} \ \& \ \neg\text{Sa})$ ' can be properly used to assert that it is not the case that Alfred drinks and does not smoke.

This is as it should be: (M1) is the correct specification of the sentence's proper communicative function. The problem is, however, that there are also uninterpretable T-sentences such as

(T2) ' $\neg(\text{Da} \ \& \ \neg\text{Sa})$ ' is true iff grass is green.

which might be true. If so, (12) and (TM) would permit derivation of a false meaning assignment such as

(M2) ' $\vdash \neg(\text{Da} \ \& \ \neg\text{Sa})$ ' can be properly used to assert that grass is green.

One might think that what is wrong with (T2) is that it is not derivable from the theory of truth for L1. Clearly, in order to know (T2) one needs to know that both its sides have the same truth value. This information cannot be derived from the theory of truth for L1. So perhaps we need to modify (TM) in the following way:

(TM*) If θ is a sentential phrase, and if ' θ is true iff p ' is a T-sentence that can be derived from the theory of truth for L1 by logic alone, then infer that θ 's content is the content that p .

(TM*) no longer applies to uninterpretable T-sentences such as (T2). However, there are still other uninterpretable T-sentences which *can* be proved from the theory of truth for L1 and logic alone. For example:

(T3) ' $\neg(\text{Da} \ \& \ \neg\text{Sa})$ ' is true iff it is not the case that Alfred drinks and does not smoke, and if grass is green then grass is green.

Again, (TM*) would allow the derivation of a false meaning-specifying theorem.

So we need to modify (TM*) again. We need to exclude T-sentences that have been derived by inferential steps that make them uninterpretable. We can do so by describing a *canonical procedure* for proving interpretable T-sentences:⁸

When deriving a T-sentence for a sentential phrase s , first apply the axiom that concerns the syntactic operation last used in constructing s to get a theorem of the form ' s is true iff ...'. For any object-language expressions mentioned in '...', repeat that procedure, and do so again for the resulting theorems until no more object-language expressions are mentioned. Then use the theorems derived, and the rule that equivalents may be substituted for one another to eliminate any object-language expressions mentioned on the right-hand side of the original theorem ' s is true iff ...'.

T-sentences derived in this way will be interpretable. So we can now use a new version of (TM):

(TM**) If θ is a sentential phrase, and if ' θ is true iff p ' is a *canonically derived* T-sentence, then infer that θ 's content is the content that p .⁹

The theory of truth and the theory of force for L1 now permit the derivation of all and only correct meaning-specifying theorems for the sentences of L1. With the addition of (TM**), the theory of truth for L1 has in effect served as a theory of content.

Let me take stock. I have shown how a theory of meaning for the very simple language L1 can be constructed by first defining a notion of sentence and sentential phrase (syntax), then defining truth for sentential phrases, and formulating an inference rule (such as (TM**)) with which one can assign to each sentential phrase a content (or proposition). Finally, one can use (12) to assign to each sentence a communicative function. Thus, one can use this theory to derive for any sentence s of L1 a meaning specification of the form ' s can be properly used to ϕ that p '. Let me call a meaning theory with this kind of architecture a 'truth-conditional meaning theory'.

L1 was just a propositional formal language, which simplified the task of formulating the definition of truth for L1. But the same could be done for languages containing quantifiers, as Tarski has shown. Now, even though Tarski's technique for defining truth for first-order formal languages is a crucial element of the truth-conditional paradigm, I do not want to clutter these pages with this

relatively standard material.¹⁰ The crucial point for the present discussion is that a truth-conditional theory of meaning for a first-order formal language is possible.¹¹

5 Application to natural languages

I have now outlined the formal structure of truth-conditional theories of meaning for formal languages. But it is far from obvious whether and how the strategy outlined for formal languages can be made to work for natural languages. There are many natural-language constructions for which it is hard to find adequate semantic axioms that would fit into a truth-conditional theory. But during the last forty years or so many philosophers and linguists have spent time trying to make recalcitrant natural-language constructions workable within the truth-conditional framework—to discover their ‘logical form’, as this project is often described. Naturally, not every problem has been treated, nor is there even unanimity about many of the issues that have been debated. Nevertheless, the range of natural-language constructions for which logical-form proposals have been formulated is impressive. Broadly conceived, the truth-conditional paradigm is now the most successful and most widely pursued approach in natural-language semantics.

The topic of this book is a difficulty that would arise even if it could be shown that there are viable analyses of the logical form of all natural-language expressions and constructions. For that reason, I will not say anything here about my preferred solutions to the perennial problems in truth-conditional semantics, such as: proper names (e.g. ‘Aristotle’), definite descriptions (e.g. ‘the King of France’), indexicals (e.g. ‘yesterday’), conditionals (i.e. if-sentences), indirect contexts (e.g. ‘Sam believes that ...’), sentences describing actions (e.g. ‘he defended himself passionately’) and so on. Rather, I shall pretend that there are satisfactory solutions to all the problems to do with the susceptibility of natural languages to formal truth-conditional treatment. For example, I shall speak of the ‘contents’ of sentences as that which is captured by the interpretive T-sentences of a truth-conditional meaning theory, without presupposing any particular view as to what kind of object a content is (if at all). Whenever details of the account are relevant for my argument, however, I will draw attention to this fact.

6 Proper communicative function

Even if we assume that questions of logical form can be settled, a lot remains to be explained, if the possibility of formulating truth-conditional theories of meaning for natural languages is to shed any light on the nature of linguistic communication via natural languages. First, I have used unexplained notions of communicative acts such as the notions of assertion and question. What is it for someone to assert or ask something? Secondly, I have not said how the meaning specifications a truth-conditional theory of meaning yields are to explain lan-

guage use. The specifications make use of the notion of proper use. What does it mean to say that a sentence ‘can be properly used’ to perform a certain communicative act?

When a theory of meaning yields a meaning specification for a given sentence type *s*, then this is supposed to license a redescription of utterances of that sentence type (by users of the language in question) as communicative (illocutionary) acts. Thus if the theory specifies ‘*s* can be properly used to assert (ask) that *p*’, then utterances of *s* may be redescribed as (literal) assertions that, or questions whether, *p*. Together with an account of the communicative acts of assertion and question, this will (eventually) allow us to explain how people can communicate by uttering and hearing sentences.

There are good reasons why it is helpful to speak of ‘proper’ use and ‘literal’ assertion. Speakers often use sentences non-literally or non-properly. The sentence ‘The door is open.’, for example, is assertoric, because it can be properly or literally used to assert that the door is open. However, it can also, in a sense, be used to issue a command. Similarly, the interrogative sentence ‘Do fish ride bicycles?’ can be used to answer a question rather than to ask one (e.g. to answer the question whether philosophers ever agree on anything). It would be very difficult to formulate a meaning theory the meaning specifications of which anticipate all possible communicative uses to which sentences may be put. The next best thing is a theory that predicts the literal or proper use to which sentences are put and leaves an explanation of non-literal uses for a separate theory. Even pre-theoretically it is not implausible to say that any utterance of the sentence ‘Do fish ride bicycles?’ is a literal question whether fish ride bicycles, even if asking that question is presumably never the ultimate point of such an utterance. Suppose that the ultimate point of one particular utterance of the sentence is to communicate that the answer to a previously asked question is obviously ‘no’. Then this aspect of the utterance should be explicable by recourse to the communicative act literally performed. This additional explanation might proceed along the lines of Grice’s theory of implicature (Grice 1989).¹²

Let me use an analogy to illustrate and justify the concentration on literal or proper use. Suppose we wanted to shed light on the function of screwdrivers, perhaps as part of a general theory of domestic tools. Suppose further that we agreed that the function of screwdrivers was somehow based on the use people make of screwdrivers. Then surely we wouldn’t be making much progress by claiming that the function of screwdrivers consisted in all the ways in which screwdrivers are actually used. Even if it were possible to specify all those ways, that would not tell us what their function is. For not all uses of screwdrivers are relevant to their function. It is the *proper uses* that are most telling. The proper uses are those uses that serve the purposes for which screwdrivers are designed, namely the purposes of screwing and unscrewing screws. We can explain the essential features of screwdrivers by recourse to these purposes (not the inessential features, such as a screwdriver’s colour). In this way, we could gain an

understanding of the function of screwdrivers and use it to explain and predict people's screwdriver-related behaviour.

Not all the improper uses of screwdrivers are wholly independent of their proper function. Some of the essential features may be exploited for purposes more or less similar to the proper purpose. I recall once using a screwdriver as a chisel. One could also use one for peeling oranges or for scratching one's back. Sometimes, it may even be unclear whether a use is proper, as for instance when I hold a bolt, instead of turning it, so that I can tighten a nut. In other cases, inessential features may be exploited in improper uses. For example, a screwdriver's colour might be used ostensively to identify a certain colour.

How does this analogy shed light on the notion of the meaning of an expression? The idea is that the meaning of a sentence consists in its proper use or proper function, and that by specifying the proper function of the sentences of a language we can predict and explain utterances of the sentence. The hypothesis is that after the proper function of sentences has been specified, we can also gain a better understanding of improper uses, much in the way in which an understanding of the proper function of screwdrivers might help in understanding improper uses of screwdrivers. Let me stress that in this connection the word 'improper' is not derogatory. There is nothing wrong in using language for purposes that diverge from proper function. There is nothing bad in using 'The door is open.' in order to issue a command, or in uttering 'Do fish swim?' not in order to ask a question. The priority of proper or literal function over improper or non-literal use is merely explanatory priority in a theory that is intended to explain the workings of a language.

So how does one decide what the proper function of a given sentence is? In the case of the screwdriver, it is easy, because the proper function of screwdrivers is determined by the intentions of the makers or designers of screwdrivers. With language, the situation is different. Although it seems to me promising to start from the assumption that language is a *tool* of communication, this tool does not have a designer or inventor (at least not in the case of natural languages). If we cannot ask 'the inventor of a language' what the proper function of its sentences is, how can we find out? My suggestion is that we should use observations of language use in order to construct a theory that explains language use. There may be many such theories, each differing in respect of which uses of language are classified as proper, but not all of these theories are equally good theories. A theory that renders all uses as proper, for example, would not be very good, because it would lack systematicity. The best theory will provide an explanation of a maximum range of uses while being simple, systematic, integrated into other theories and using pre-theoretically familiar concepts (compare Thagard 1978).

On some such theory, I think, deviant uses will no longer be a problem. They will either be so deviant that they do not fall into the range of uses covered by the theory, or they will occur as non-basic uses that can be explained by recourse to the proper, basic function of the expressions used.

On this view, the literal meaning of a sentence, or its proper communicative function, therefore enjoys theoretical priority over the specific non-literal uses to which it may be put by language users. Is this methodologically justified? After all, a theory of meaning for a particular language should be based on the use people make of that language. The sentence 'Do fish ride bicycles?' has presumably never been used, in the pre-theoretical sense, to ask whether fish ride bicycles. So with what justification can one say that it can be *properly* used to ask whether fish ride bicycles? How do the data of language use confirm that the sentence has this communicative function if no one uses it to fulfil that function? Can reasons of theoretical simplicity justify this apparently blatant disregard for (or perhaps even inconsistency with) the data?

The answer is that the sentence's force indicators and content indicators are used as part of other sentences in ways that confirm that this is the sentence's proper function. Its word order and introduction with the auxiliary verb 'to do', as well as the question mark (or a certain inflection in spoken English), are syntactic devices that are used in countless other sentences that do in fact get used to ask questions—for example the sentence 'Does Sam smoke?'. The prediction that the sentence's proper function is interrogative is based on these data. The sentence's content indicators 'fish', 'ride', 'bicycle' and the way in which they are compounded are also devices that display their proper function in other sentences—for example 'Dutchmen ride bicycles.' or 'Fish have scales.'

The systematicity at which a theory of meaning of this kind aims should therefore reflect the systematic, repeatable role each syntactic device displays in many different contexts. This methodological maxim serves an explanatory aim of a meaning theory. As I have mentioned earlier, one of the aspects of language use that needs explanation is the ability of language users to use novel sentences. If a theory of meaning for a language systematically describes the repeatable contribution each syntactic part makes to the proper function of the sentences in which it may feature, then such a theory can serve as a model for the cognitive processes by which language users compute the proper function of sentences of their language. Thus the division into uses that manifest a sentence's proper communicative function directly and those for the explanation of which we need to appeal to additional principles governing improper uses (e.g. a theory of implicature, ellipsis or metaphor) is a distinction that should be anchored in the actual linguistic behaviour of speakers.

Let me, once again, take stock. A truth-conditional theory of meaning for a language will yield, via axioms concerning all the content indicators and force indicators of the language, meaning specifications for all the sentences. These meaning specifications take the form '*s* can be properly used to ϕ that *p*'. These specifications in turn license the redescription of utterances as (proper or literal) communicative acts, e.g. 'A has (literally) asserted that *p*'. Thus someone who possessed the theory could derive, for each utterance of a sentence of the language, which (literal) communicative act it serves to perform.¹³ In this sense he or she would possess the ability to understand the language.

7 Communicative acts

Finally, I need to explicate the notion of a communicative act itself, i.e. the notions of assertion and question and perhaps further types of communicative act if necessary. These notions are no doubt pre-theoretically familiar. But some explication of the notions is needed to make their theoretical role clear. In §6, I said that a theory of meaning for a natural language can be empirically confirmed by data of language use. In that connection I talked as if we could observe directly that some sentence is used on some occasion to perform a certain communicative act. Thus a speaker's uttering a sentence *s* and thereby asserting that *p* confirms, to some degree, any theory of meaning that entails that *s* can be properly used to assert that *p*. Obviously, this can work only if there is some way of knowing that the speaker asserted that *p* by uttering *s* independently of having a correct meaning theory for the language in question. What we need is an explication of the notion of assertion, and of other kinds of communicative act, in independent terms.

In the literature, this problem is often discussed in the guise of the quest for an adequate characterization of the 'actual language relation'.¹⁴ A theory of meaning of the kind I have outlined can be viewed as the description of an abstract object—a complex function that assigns meanings to sentence types. So conceived, there are countless theories of meaning that describe countless ('possible') languages. Only very few of these possible languages are actual languages, i.e. languages used by some population of language users.¹⁵ The actual language relation is the relation in which a possible language stands to a population when, and only when, that population uses that language. What does this relation consist in?

I shall not here address the problem of the actual language relation directly, but rather approach the problem from the angle already mentioned, i.e. by looking for an account of the various communicative acts. I shall briefly outline an account of the sort I favour.

What we need is an account of assertion that makes it possible to judge whether it is likely, given general psychological assumptions, that an agent is intending to assert a certain content when uttering a given sentence type. One account of assertion (Brandom 1983) describes the linguistic practice of making assertions as a kind of game. In any game, there is a number of possible moves, and each move has a significance within the game. According to Brandom, to assert is to make a move in the social game of communication, a game the rules of which are ultimately underpinned by social norms and sanctions. The significance of making an assertion can be described as follows: to assert that *p* is (1) to undertake a responsibility to justify *p* if challenged to do so, and (2) to issue a licence to use *p* as a premise. Thus if one participant makes an assertion, other participants may challenge it. Since the asserter has, by making the assertion, undertaken a responsibility to justify his or her assertion if challenged, he or she is then obliged to justify the assertion. This justification can take the form of a

further assertion which, if true, would confirm the original assertion, or the asserter may invoke an assertion made by another participant. If the asserter fails to comply with this obligation, he or she will normally suffer negative sanctions (e.g. his or her status in the community may suffer in some way). Since an asserter is under an obligation to provide justifications, other participants will frequently rely on an assertion and take it as evidence that what has been asserted is true—especially if they believe that the asserter would not have made the assertion, and risked sanction, unless he or she was in possession of good justification). Thus agents who make assertions will often bring it about that their addressees acquire new beliefs or at least change the likelihood they assign to what has been asserted. This is how participants in the game of assertion use language to communicate.

If this account is correct, then it provides a way of using general psychological principles to judge whether a given population with the patterns of its speech behaviour is in fact using a certain language. Suppose a participant utters a sentence the utterance of which counts as an assertion that *p* in a language *L*, and that general psychology renders it likely that he or she would want to undertake the responsibilities and issue the licences that making the assertion involves. Then the view that members of that population use *L* would be confirmed to some degree. Similarly, the participants' responses to assertions can provide evidence as to whether the population uses *L*: if the assertion is followed by what counts as a challenge in *L*, and this is then followed by what counts as a justification in *L*—or by a failure to justify and ensuing sanction—then this equally confirms the hypothesis that the population uses *L*.

An account of the communicative act of questioning could follow similar lines. By asking a question, a participant imposes an obligation on his or her addressee to answer the question or else give an excuse for failing to answer the question. More precisely, to answer a question whether *p* is to assert either *p* or the negation of *p*. A typical excuse for failing to answer a question is an assertion that one doesn't know whether *p*. Again, if this account of the communicative act of asking a (yes/no) question is correct, then it provides a basis for judging whether a given utterance is likely to have been intended as a question.

Of course this is only the barest outline of an account. Our actual communicative practices are far more complicated and allow for many exceptions. For example, under certain conditions, the obligation to answer a question may be overruled by other considerations, e.g. of secrecy. Similarly, some assertions are in no further need of justification, e.g. because their content is taken to be obviously true, or because the asserter has special authority or expertise. Moreover, the rules of the game do not only serve the aim of spreading information. Occasionally, participants will lie, or otherwise rely on the rules of the game, in order deliberately to mislead or deceive other participants. In those cases, participants run an (often calculated) risk of discovery and subsequent negative sanctions.

The general idea, however, is clear: communicative acts such as asserting and questioning have a certain social significance among users of the language and