

Introduction

It seems reasonable to suppose that the expressions we use as speakers of a language such as English mean something. It is also natural to assume that, under appropriate conditions, the employment of at least some among these expressions achieves effects describable with the help of locutions such as 'true' or 'false'. Finally, it is at least *prima facie* sensible to hypothesize that there is an interesting connection between these dimensions, and that a philosophically interesting story may be told about the relationship between meaning, truth, and the use of language. An important research programme within linguistics and philosophy of language, sometimes called 'natural language semantics' or 'formal semantics', is grounded on a particular notion of how such a story, or at least an important portion of it, is supposed to go. The aim of the present book is to clarify the understanding of meaning and truth that lies at the basis of the aforementioned programme, to explain how it may be applied to particular instances involving the use of language, and to defend it against an increasingly fashionable sceptical attitude. The projects of clarification and defence are complementary. The critics of the traditional paradigm, so I argue, proceed from incorrect assumptions about its scope and structure. Still, their

mistake is understandable: to an important extent, the traditional approach to semantics has been misunderstood even by its foremost defenders, in particular when it comes to the theory of meaning and truth upon which it is grounded.

1. The Plan: Chapters 1–3


The traditional paradigm within formal semantics has generated a multiplicity of different proposals, focused on alternative features of the semantic behaviour of natural languages. In this book I focus on the type of structures that emerged a few decades ago within the debate on so-called indexical languages. Although a variety of alternative approaches would do for my purpose, I concentrate for concreteness' sake on treatments somewhat reminiscent of those developed within the Montaguevian tradition, and now typically associated with the work of Hans Kamp, David Kaplan, and David Lewis. As I explain in Chapter 1, formal approaches of this type are specifically interested in certain aspects of contextual dependence: namely, those relevant for the interpretation and evaluation of indexical expressions. Simple indexical expressions, such as 'I' or 'now', refer to distinct items with respect to alternative parameters, say, the person who is speaking or the time of utterance, and they apparently do so in virtue of certain aspects of their conventional meaning. For this reason, the study of languages of this ilk provides a particularly fertile ground for the discussion of the interface between questions of meaning, issues of reference and truth, and at least certain forms of the contextual sensitivity uncontroversially affecting our linguistic interchanges. In particular, according to the classical view, the analysis of this interface reveals important logical properties of certain expressions; that is, it uncovers constructions which, in the traditional parlance, are 'true in virtue of meaning'.

2 ~ Introduction

Chapter 1 is devoted to a general and relatively informal explanation of the structures traditionally employed for the analysis of simple linguistic fragments involving indexical expressions. These structures take certain abstract items as input, and yield assignments of truth-values and, consequently, of logical properties and relations. As will emerge later in this essay, widespread mistakes regarding the scope and function of such structures may in part be traced to the formally unobjectionable, but pedagogically misleading choice of certain labels for a variety of the aforementioned parameters. For instance, it is common to refer to the *analysanda* in the project under discussion as ‘sentence–context pairs’, and to the results with which they are paired as ‘truth-conditions’. It is also customary to label the theoretical machinery designed for the assignment of truth-conditions to such pairs as a procedure of ‘semantic’ compositional interpretation. In order not to prejudge a variety of issues that eventually become of immediate concern in later chapters, I opt for an artificial, deliberately neutral terminology: I thus talk of *interpretive systems* (or, more often, simply *systems*) which, when applied to *clause–index* pairs, yield conclusions of *t-distributions*, i.e., assignments of truth-value at particular *points* of evaluation.

One of the didactic pay-offs for this unwieldy terminology consists in the rather obvious chasm it imposes between the interpretive system’s concerns, on the one hand, and, on the other, the questions pertaining to its *application* to the nitty-gritty of everyday language use. Systems assign *t-distributions* to clause–index pairs, but competent and intelligent speakers are attuned to more tangible dimensions: in particular, to the intuitive truth-values of particular utterances on given occasions. If systems may eventually come to grips with such intuitions and aim at results consonant with them, they may do so only on the assumption of appropriate hypotheses about their interface with the world of daily exchanges—hypotheses pertaining to the clause–index pair adequate for the representation of an utterance,

and to the understanding of the system's t-distributional output in truth-conditional terms. In a more pictorial fashion



The discussion of the 'gaps' between, on the one hand, the system's input (a clause-index pair) and output (a t-distribution), and, on the other hand, the intuitive parameters of semantic analysis (an utterance's truth-conditions), is one of this book's main concerns. In the final sections of Chapter 1, I begin to address the relationship between an utterance—that is, an instance of language use taking place in a given context—and the clause-index pair appropriate for its analysis. Armed with the discussion of such relationship, I critically approach some considerations put forth by the defenders of a fashionable sceptical standpoint having to do with issues such as disambiguation or reference assignment.

In Chapter 2, I continue my discussion of how utterances may be appropriately represented from the interpretive system's point of view. In particular, I focus on the relationship between the *context* in which an utterance takes place and the index involved in its representation. The starting-point for this discussion is provided by rather frivolous cases, having to do with recorded messages and written notes. But the point which these examples help to uncover transcends the not-so-urgent need for a theory of postcards or answering machines. The main conclusion of this chapter is that even some of the foremost defenders of the customary treatment of indexical languages have burdened traditional systems with extraneous assumptions, thereby concealing the view of meaning and truth to which they are truly committed. The methodological gains of my non-traditional labels, in particular my cautious distinction between contexts and indexes, are apparent in this respect. It is indeed advisable, at least at a preliminary stage, that questions related to the parameters

selected by the meanings of indexicals be isolated from the assumption that what is being addressed is a context, in the everyday sense of the concrete setting in which an utterance takes place. The relationship between a context and what I call an ‘index’, namely the collection of items requested by the meaning of the indexicals under analysis, is non-trivial, and should not be prejudged by unwarranted terminological decisions.

Chapter 3 continues the discussion of the relationships between the interpretive system’s mechanisms and its application to particular utterances. In agreement with the founding fathers of the traditional treatment of indexical languages, systems of the type sketched in Chapter 1 operate on clause–index pairs; i.e., they evaluate expression-types with respect to appropriate additional parameters. It has, however, occasionally been suggested that an alternative, and possibly more appropriate, approach to indexicality eschews expression-types, in favour of a *reflexive* analysis geared towards their particular exemplars. The motivation behind this suggestion, or at least behind the versions of so-called utterance semantics in which I am interested, is semantic, rather than metaphysical. What is at issue is not the antipathy for abstract instantiables, such as, presumably, types, or the predilection for the everyday concreteness of tokens. The point has rather to do with the aims and scope of a systematic analysis able to yield results for utterances, i.e., with the aims and scope of ‘applied interpretive systems’. It is on these terms that I take up the utterance-semanticist’s challenge. The conclusion I reach is negative: on an appropriate understanding of meaning and truth, interpretive systems had better steer clear of the structures proposed by token-reflexive approaches.

2. *The Plan: Chapters 4–5*

Chapters 2 and 3 pursue different themes related to the appropriate *input* for an interpretive system, and to its applications to

particular utterances. The discussion of the relationship between clause–index pairs and utterances is important, because systems—namely, procedures that operate on the former—aim at empirical adequacy; i.e., at consistency with pre-theoretic intuitions pertaining to the latter. What is desired, among other things, is that the interpretive system, when supplied a clause–index pair appropriate to a certain utterance *u*, gives results suitably related to (at least some among) our intuitive verdicts about *u*. But the interface between the system’s theory of meaning and truth, on the one hand, and the treatment of particular instances of language use, on the other, does not only raise questions pertaining to the input on which the former operates. As highlighted by the deliberately artificial terminology I adopt in Chapter 1, what systems yield are results of *t-distributions*. Yet, what our intuitive assessments puts forth are not judgements of truth-values at particular points of evaluation, but conclusions of *truth-conditions*. It is to the analysis of the relationship between *t-distributions* and *truth-conditions*—that is, in the figurative jargon introduced above, to the discussion of the second ‘gap’ separating interpretive systems from everyday intuitions—that Chapter 4 is devoted.

It is here that I return to the fashionable contextualist attacks on traditional structures that I began addressing in Chapter 1. Leaving aside the additional worries briefly addressed there, having to do with reference assignment or ambiguity resolution, the contextualists ground their challenge on the conviction that customary treatments of meaning and truth are empirically inadequate: the view of meaning and truth presupposed in Chapter 1, so it is claimed, often yields incorrect conclusions of *truth-conditions*. I disagree: once the aim and scope of a traditional interpretive system are properly understood, the intuitively required *truth-conditional* outcomes are perfectly consistent with that system’s *t-distributional* results.

As this preliminary summary of the first four chapters indicates, this essay’s main concern is of a ‘metasemantic’ nature:

what I address are the philosophical and theoretical commitments of treatments of a particular type, in particular their commitment to certain views about meaning and truth. Yet, the discussion of these general issues is of relevance not only from the point of view of the assessment of an influential research programme, but also for a variety of questions ‘internal’ to it. If my considerations in the first four chapters of this essay are correct, the traditional approach to meaning and truth has been misunderstood, to varying extents, not only by the foremost contextualist sceptics, but also by many who take a friendlier attitude towards it. Unsurprisingly, this misunderstanding has affected the treatment of a multitude of problems that typically occupy traditional semanticists in their everyday toil, regardless of their reactions to one or another among the challenges to the core assumptions within their paradigm. I attempt to substantiate this contention by example. One will have to suffice, but this lack is at least partially compensated by the fact that what I confront in the final chapter of this essay is one of the most discussed semantic problems of recent decades: the treatment of attitude reports and of singular terms occurring within them. Chapter 5 argues for the conclusion that, once the aims and structure of the interpretive system are properly understood, the problem raised by occurrences of singular terms within attitude reports is an ‘easy’ one, in the sense that it does not require the negation of any among the most straightforward views regarding reference, attitude predicates, complementizers, and the like.