

Preface

This book has three aims. The first is to motivate and explain a class of approaches to the semantics of natural languages known as *dynamic semantics*. The second is to give an extended criticism of those approaches. The third is to present and develop an alternative theory that better accounts for the kinds of patterns that motivate dynamic semantics.

Why bother with an extended exposition of dynamic semantics if only to argue against it? The reason is that there is a lot that is insightful about dynamic semantics. Dynamic semantics is right to seek a uniform treatment of patterns involving superficially unrelated constructions, including modality, conditionals, and anaphora. And it is right to explain the parallels across those domains using the notion of *local information*: the information available to a part of a sentence, given the globally available information and the rest of the sentence.

Nonetheless, dynamic semantics fails in a variety of related ways. To get a sense for this, consider briefly the dynamic treatment of conjunction. Dynamic semantics treats sentence meanings, not as functions from a context and index to a truth-value—as on traditional (“static”) approaches—but rather as *context change potentials (CCPs)*: functions which take an input context to an output context. Corresponding to this is a non-classical treatment of sentential connectives. Conjunction, in particular, is standardly treated as successive updating: the meaning of $p \wedge q$ first applies the meaning of p to a context, and then applies the meaning of q to the resulting context (symbolically, where $[\cdot]$ takes a sentence to the corresponding CCP, $c[p \wedge q] = [q]([p](c))$ for any context c). This means that the context that the meaning of q operates on is different from the one that the meaning of p operates on.

This kind of intra-sentential dynamic is the heart of dynamic semantics, and central to its signal empirical successes. But it also leads to pervasive problems. Among these are logical problems. To take one example, the Law of Non-Contradiction, which says that $p \wedge \neg p$ is never true, fails. Abstractly, the point is this: $p \wedge \neg p$ denotes a successive update which takes a context c , applies $[p]$ to it, and then applies $[\neg p]$ to the result. Since $[\neg p]$ thus updates a different context from $[p]$, it is, in a loose sense, interpreted differently from $[p]$. That means that there is no reasonable meaning for negation in this framework which guarantees that the result of this succession of updates is contradictory.

Non-classicality, of course, is not an objection per se. One of the points of dynamic

semantics, after all, is to capture ways that natural language diverges from classical logic. But, as I will argue, many of the particular ways that dynamic semantics diverges from classical logic do not match intuitions about inference and coherence.

Dynamic semantics also has related problems with order. One of the key ideas of dynamic semantics is that conjunction is a lot like successive assertion. This allows dynamic semantics to capture key similarities in the interpretation of *sequences* of assertions like ‘There is a cat in the tower. He is sulking’, and corresponding *conjunctions* like ‘There is a cat in the tower, and he is sulking’. The treatment of conjunction as successive update encodes order-sensitivity in the meaning of conjunction: conjunction updates first with the left conjunct, then the second, so that $p \wedge q$ and $q \wedge p$ will not in general mean the same thing, unlike on the classical static approach. But I will argue that many of the patterns which dynamic semantics aims to capture are in fact order-insensitive. This is a problem for the dynamic idea that conjunction is essentially a semanticization of successive assertion.

These problems, I will argue, arise from the way that dynamic semantics ties together the dynamics of local information and truth-conditions into a single dimension. The central idea of the alternate approach which I will develop, the *bounded theory*, is that, while local information influences interpretation, that influence is indirect, via a dimension of meaning distinct from truth-conditions. On the bounded theory, at the level of truth-conditions, natural language has a broadly classical, static semantics. The dynamic influence of local information comes via a separate dimension of constraints—a dimension I call *bounds*—which limit admissible interpretations of context-sensitive language. This approach captures the characteristic patterns that motivate dynamic semantics, and it does so without implausible predictions about logic and order.

The problems for dynamic semantics which I will explore differ in detail. But there are striking parallels in the failures, and there is a common move available in response to them: locating the characteristically dynamic features of interpretation in a separate dimension of bounds, while leaving the truth-conditions static. What emerges is a new perspective on how meaning is, and is not, dynamic.

What emerges is also a collection of detailed proposals about the truth-conditions and bounds of different fragments. These will, I expect, be improved upon in future work; and I hope that work will find useful ingredients in both the general picture and the particular proposals which I develop. More broadly, I hope this book will facilitate progress by introducing researchers to dynamic phenomena in natural language and to a variety of treatments of them (dynamic semantics, the bounded theory, and a few other approaches we’ll explore along the way). My topic lies at the intersection of linguistics and philosophy, with connections to logic and sometimes computer science. That is both exciting and challenging. Bringing together an overview of the motivating linguistic phenomena, the basic ideas of the dynamic approach, and an exploration of alternatives will make the area accessible across a range of disciplines, and, I hope, open up new avenues of research beyond the horizon of the present book.