The Cambridge Handbook of

Formal Semantics

edited by **Maria Aloni**
and Paul Dekker

The Cambridge Handbook of Formal Semantics

Formal semantics – the scientific study of meaning in natural language – is one of the most fundamental and longest-established areas of linguistics. This handbook offers a comprehensive yet compact guide to the field, bringing together research from a wide range of world-leading experts. Chapters include coverage of the historical context and foundation of contemporary formal semantics, a survey of the variety of formal/logical approaches to linguistic meaning, and an overview of the major areas of research within current semantic theory, broadly conceived. The handbook also explores the interfaces between semantics and neighbouring disciplines, including research in cognition and computation.

This work will be essential reading for students and researchers working in linguistics, philosophy, psychology, and computer science.

MARIA ALONI and PAUL DEKKER are Assistant Professors at the Institute for Logic, Language and Computation (ILLC) of the University of Amsterdam. Both are leading scholars and researchers in the area of logical semantics, the philosophy of language, and empirical linguistics. Their general interests include the formal semantics and pragmatics of natural language, and more specific subjects such as reference, quantification, indefiniteness, modality, and information structure, and general issues on the syntax–semantics and the semantics–pragmatics interface.

CAMBRIDGE HANDBOOKS IN LANGUAGE AND LINGUISTICS

Genuinely broad in scope, each handbook in this series provides a complete state-of-the-field overview of a major sub-discipline within language study and research. Grouped into broad thematic areas, the chapters in each volume encompass the most important issues and topics within each subject, offering a coherent picture of the latest theories and findings. Together, the volumes will build into an integrated overview of the discipline in its entirety.

Published titles

The Cambridge Handbook of Phonology, edited by Paul de Lacy
The Cambridge Handbook of Linguistic Code-switching, edited by Barbara E. Bullock and Almeida Jacqueline Toribio
The Cambridge Handbook of Child Language, Second Edition, edited by Edith L. Bavin and Letitia Naigles
The Cambridge Handbook of Endangered Languages, edited by Peter K. Austin and Julia Sallabank
The Cambridge Handbook of Sociolinguistics, edited by Rajend Mesthrie
The Cambridge Handbook of Pragmatics, edited by Keith Allan and Kasia M. Jaszczolt
The Cambridge Handbook of Language Policy, edited by Bernard Spolsky
The Cambridge Handbook of Second Language Acquisition, edited by Julia Herschensohn and Martha Young-Scholten
The Cambridge Handbook of Biolinguistics, edited by Cedric Boeckx and Kleanthes K. Grohmann
The Cambridge Handbook of Generative Syntax, edited by Marcel den Dikken
The Cambridge Handbook of Communication Disorders, edited by Louise Cummings
The Cambridge Handbook of Stylistics, edited by Peter Stockwell and Sara Whiteley
The Cambridge Handbook of Linguistic Anthropology, edited by N.J. Enfield, Paul Kockelman and Jack Sidnell
The Cambridge Handbook of English Corpus Linguistics, edited by Douglas Biber and Randi Reppen
The Cambridge Handbook of Bilingual Processing, edited by John W. Schwieter
The Cambridge Handbook of Learner Corpus Research, edited by Sylviane Granger, Gaëtanelle Gilquin and Fanny Meunier
The Cambridge Handbook of English Historical Linguistics, edited by Merja Kytö and Päivi Pahta
The Cambridge Handbook of Linguistic Multi-competence, edited by Li Wei and Vivian Cook

Forthcoming

The Cambridge Handbook of Historical Syntax, edited by Adam Ledgeway and Ian Roberts
The Cambridge Handbook of Linguistic Typology, edited by Alexandra Y. Aikhenvald and R. M. W. Dixon
The Cambridge Handbook of Morphology, edited by Andrew Hippisley and Greg Stump

The Cambridge Handbook of Formal Semantics

Edited by
Maria Aloni and **Paul Dekker**
University of Amsterdam

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781107028395

First published 2016

Printed in the United Kingdom by Clays, St Ives plc

A catalogue record for this publication is available from the British Library

ISBN 978-1-107-02839-5 Hardback

Contents

Figures

Tables

Contributors

Maria Aloni, Assistant Professor of Logic and Language, University of Amsterdam
Nicholas Asher, Professor of Linguistics and Philosophy, Université Paul Sabatier
Giosuè Baggio, Associate Professor, Norwegian University of Science and Technology
Adrian Brasoveanu, Associate Professor of Linguistics, University of California at Santa Cruz
Lucas Champollion, Assistant Professor of Linguistics, New York University
Ariel Cohen, Associate Professor of Linguistics, Ben-Gurion University of the Negev
Mikaël Cozic, Assistant Professor of Philosophy of Science, Université Paris-Est and Institut Universitaire de France
Paul Dekker, Assistant Professor of Logic and Language, University of Amsterdam
Paul Egré, CNRS Research Fellow, Institut Jean-Nicod; Associate Professor of Logic and Philosophy, Ecole Normale Supérieure
Donka Farkas, Professor of Linguistics, University of California at Santa Cruz
Jonathan Ginzburg, Professor of Linguistics, Université Paris-Diderot
Jeroen Groenendijk, Professor of Logic and Language, University of Amsterdam
Atle Grønn, Professor of Russian Linguistics, University of Oslo
Hans Kamp, Professor of Formal Logics and Philosophy of Language, Universität Stuttgart
Manfred Krifka, Professor of General Linguistics, Humboldt-Universität zu Berlin
Michiel van Lambalgen, Professor of Logic and Cognition, University of Amsterdam
Lisa Matthewson, Professor of Linguistics, University of British Columbia

Louise McNally, Professor of Linguistics, Universitat Pompeu Fabra
Rick Nouwen, Associate Professor of Linguistics, Utrecht University
Peter Pagin, Professor of Theoretical Philosophy, Stockholm University
Barbara Partee, Distinguished University Professor Emerita of Linguistics and Philosophy, University of Massachusetts, Amherst
Paul Portner, Professor of Linguistics, Georgetown University
James Pustejovsky, Professor of Computer Science, Brandeis University
Susan Rothstein, Professor of Linguistics, Bar-Ilan University
Manfred Sailer, Professor of Linguistics, Johann Wolfgang Goethe-Universität Frankfurt
Galit W. Sassoon, Senior Lecturer in Language and Cognition, Bar-Ilan University
Philippe Schlenker, Director Research, Institut Jean-Nicod, CNRS; Global Distinguished Professor of Linguistics, New York University
Arnim von Stechow, Professor of General and Theoretical Linguistics, Eberhard Karls Universität Tübingen
Keith Stenning, Professor of Informatics, University of Edinburgh
Matthew Stone, Professor of Computer Science, Rutgers University
Henriëtte de Swart, Professor of French Linguistics and Semantics, Utrecht University
Enric Vallduví, Associate Professor in Translation and Language Science, Universitat Pompeu Fabra
Dag Westerståhl, Professor of Theoretical Philosophy and Logic, Stockholm University
Thomas Ede Zimmermann, Professor of Formal Semantics, Johann Wolfgang Goethe-Universität Frankfurt

Preface

Humans characteristically use language, and a characteristic feature of the use of language is that it is meaningful. Semantics is the study of meaning, of the structural ways in which it is realized in natural language, and of the formal logical properties of these structures. The area of formal semantics finds its roots in logic and the philosophy of language and mind, but it has also become deeply entrenched in linguistics and the cognitive sciences.

This *Cambridge Handbook of Formal Semantics* constitutes a comprehensive guide to contemporary formal semantics, and it provides, among other things, a historical context and foundation of the field, a survey of the variety of formal/logical approaches to linguistic meaning, an overview of the major areas of research within current semantic theory, and a presentation of the interfaces between semantics and other domains of linguistic inquiry, broadly conceived.

This handbook is intended for everyone interested in the understanding of meaning. It presents a broad view of the semantics and logic of natural language and, as a helpful tool, of the logical languages employed.

The twenty-five chapters constituting this handbook have been grouped together into five major parts, and we hope the handbook can thus be seen to cover, in a systematic and transparent way, both the broad and the varied scope of the domain as well as the width and variety of the perspectives adopted. The contributions are subsumed under the following headings:

I The landscape of formal semantics.
II Theory of reference and quantification.
III Temporal and aspectual ontology and other semantic structures.
IV Intensionality and force.
V The interfaces.

The first and the last parts are more of a methodological or programmatic nature. The first part gives a general sketch of the frameworks in and perspectives from which semantic research is conducted. The last part focuses on the intimate and intrinsic relationships with bordering academic disciplines. The three middle parts are concerned with more or less established major domains of linguistic research: the nominal domain (Part II) and the verbal or predicative domain (Part III). Both domains are studied with logical and (natural language) ontological interests. Part IV gives an overview of the various moods and modalities in language, ranging from negation, to modals, questions and other moods.

All contributors to the handbook are senior researchers working in logic and linguistics who are widely recognised as experts on the topics assigned to them. All contributions have been thoroughly reviewed by expert colleagues.

This handbook is aimed at graduate and PhD students, but it is also meant for a wider audience, and hopes to serve as a manual and source book for senior lecturers and researchers alike. All five parts center around formal semantic issues but are intended to be of interest to linguists in general, descriptive semanticists, philosophers of language, and those who work on dialogue systems, cognitive grammar, and computation.

This work is intended to document the background and development of the main currents in formal semantics. It is not intended as a textbook, although the individual chapters may figure well in introductory and advanced courses, supplying required overviews and suggestions for further reading. For more detailed introductions to the various frameworks and formalisms, we are confident that readers can find their way among a wealth of standard textbooks. (The following list is incomplete.) G. Chierchia and S. McConnell-Ginet, *Meaning and Grammar* (1990), L. T. F. Gamut, *Logic, Language and Meaning* (1991), A. Kratzer and I. Heim, *Semantics in Generative Grammar* (1998), S. Löbner, *Understanding Semantics* (2002), B. Partee, A. ter Meulen, and R. Wall, *Mathematical Methods in Linguistics* (1990), W. Sternefeld and T. E. Zimmermann, *Introduction to Semantics* (2013), Y. Winter, *Elements of Formal Semantics* (2015).

A handbook like this would of course not have been possible without the help of the publisher, and we would like to thank especially Helen Barton for her continuous support during the lengthy process. Peter van Ormondt has proven extremely valuable as the technical editor, who, it must be emphasized, also contributed his valuable expertise on substantial matters. All of this is greatly appreciated. We are also very grateful to the reviewers of the individual chapters, whom we would have liked to list if the reviews had not been done anonymously. Their acute and timely comments have surely increased the value of the various contributions even more. Lastly, we would like to thank the individual and collective authors, for enabling us to make this broad subject, with the help of their expertise, views, and findings, accessible to a wide audience.

Maria Aloni
Paul Dekker
Amsterdam

Part I

The landscape of formal semantics

1

Formal semantics

Barbara H. Partee

1.1 Formal semantics: what is it?

Formal semantics is an approach to semantics, the study of meaning, with roots in logic, the philosophy of language, and linguistics. The word *formal* in "formal semantics" is opposed to *informal* and reflects the influence of logic and mathematics in the rise of scientific approaches to philosophy and to linguistics in the twentieth century. Distinctive characteristics of this approach (see also Pagin, Chapter 3) have been truth conditions as a central part of meaning; (usually) a model-theoretic conception of semantics; and the methodological centrality of the Principle of Compositionality: "The meaning of a whole is a function of the meanings of its parts and their mode of syntactic combination." Most formal semantics is model-theoretic, relating linguistic expressions to model-theoretically constructed semantic values cast in terms of truth, reference, and possible worlds or situations (hence, formal semantics is *not* "formal" in the sense of Hilbert, 1922). And most formal semanticists treat meaning as mind-independent (and abstract), not as concepts "in the head"; formal semanticists distinguish semantics from knowledge of semantics (Lewis, 1975b; Cresswell, 1978).

This sets formal semantics apart from approaches which view semantics as relating a sentence principally to a representation on another linguistic "level" (*logical form*) (May, 1985) or a representation in an innate "language of thought" (Fodor, 1975) or "conceptual representation" (Jackendoff, 1992). The formal semanticist could accept such representations as an aspect of semantics but would insist on asking what the model-theoretic semantic interpretation of the given representation-language is (Lewis, 1970). Kamp's Discourse Representation Theory is an exception, since as noted in Section 1.3.3 below, it includes as essential an intermediate level of representation with claimed psychological reality. Formal semantics is centrally concerned with *compositionality* at the syntax–semantics interface (see Sailer, Chapter 21), how the meanings of larger constituents are built up from the meanings of their parts on the basis of their syntactic structure.

The most important figure in the history of formal semantics was undoubtedly Richard Montague (1930–1971), whose seminal works in this area date from the late 1960s and the beginning of the 1970s. Other important contributors will also be discussed below. Since the 1980s formal semantics has been a core area of linguistic theory; important contributions also continue to come from philosophy, logic, cognitive science, and computational linguistics.

In the last thirty years formal semanticists have become increasingly concerned with issues at the interface between semantics and pragmatics, including context-dependence, information structure, and the semantics/pragmatics of dialogue (see Asher, Chapter 4; Ginzburg, Chapter 5; Schlenker, Chapter 22; Vallduví, Chapter 23). These broadening concerns have led to a range of newer approaches that treat meaning as something more than truth conditions, but still including truth conditions, possibly derivatively, as a central part of what semantics is to capture.

In this chapter we briefly trace the history of formal semantics (Section 1.2) and discuss some of its central principles, debated issues, and divergences within the field (Section 1.3). Since issues concerning the syntax–semantics interface are so crucial to the central working hypothesis of compositionality, we include some brief case studies relating to the syntax–semantics interface in Section 1.4; fuller treatments of related issues will be found in Brasoveanu and Farkas, Chapter 8, and Sailer, Chapter 21. In Section 1.5 we describe the increasing attention to the role of context and to language use and the consequent blending of formal semantics and formal pragmatics (see also Schlenker, Chapter 22). In Section 1.6 we come back to the foundational question of whether meanings are in the head and how formal semantics, which has traditionally rested on the assumption that they are not, connects to cognitive science and studies of language processing and language acquisition. In the final Section 1.7, we mention some of the relatively recent contributions of formal semanticists to issues of language universals and language typology.

1.2 The history of formal semantics

1.2.1 Semantics in linguistics before 1970

Linguistics had partly different beginnings and different emphases in Europe and in America, growing in considerable part out of philological-historical work in Europe and out of that plus anthropological studies in America. Semantics in early linguistics was mainly lexical; lexical semantics and principles of semantic change and semantic drift were important for historical and comparative linguistics. Structuralism arose first in Europe, and de Saussure was influential for structuralism, for putting synchronic grammar into the foreground, and for conceiving of grammar as connecting form and meaning. Bühler's *Sprachtheorie* (1934) included an early treatment of indexicality and perspective shift. Jespersen made lasting contributions to semantics as well as syntax (1924); while in the Netherlands, Evert Beth was laying foundations (1947, 1963) for the cooperation among logicians and linguists that has made the Netherlands one of the major contributors to the development of formal semantics.

Semantics was rather neglected in early and mid-twentieth-century American linguistics, for several reasons. Early American anthropological linguistics depended on fieldwork, where phonetics came first, then phonology, morphology, perhaps a little syntax, and usually no semantics beyond a dictionary. Another reason was the influence of logical positivism and behaviorism: meaning was viewed as an unobservable aspect of language, not fit for scientific study. And the Chomskyan revolution concentrated on syntax: there was much talk of the creativity of language and of language as a window on the mind, but it was all about syntax, investigating finite recursive mechanisms that could characterize the infinite class of sentences of a natural language, and trying to solve the mystery of first-language acquisition.

In 1954, the philosopher Yehoshua Bar-Hillel wrote an article in *Language* (1954) urging cooperation between linguists and logicians and arguing that advances in both fields made the time ripe for combining forces to work on syntax and semantics together. However, Chomsky immediately replied (1955), arguing that the artificial languages invented by logicians were too unlike natural languages for any methods the logicians had developed to have any chance of being useful for developing linguistic theory. Real cooperation of the sort that Bar-Hillel urged began only after Montague's work.

The first efforts to add semantics to Chomsky's generative grammar were made by Katz and Fodor (1963), who were the first to use the term *compositionality* (although in most of their work they speak rather of *projection rules*), proposing rules that first interpret underlying phrase-markers, bottom-up, and then interpret the result of applying transformations. A little later Katz and Postal (1964) proposed that all semantically relevant information should actually be contained in the underlying phrase-markers and that transformations should preserve meaning. The Katz–Postal hypothesis was adopted in Chomsky (1965) as part of what Chomsky later dubbed "the standard theory," with Deep Structure as the input to semantics.

But very soon the "discovery" of quantifiers spoiled the illusion of meaning-preservingness for many transformations; I illustrate with a few key examples without giving the rules or the structures. In each case the (a) example illustrates the apparently meaning-preserving transformation when the NPs are proper names, but in the (b) example, application of the same rule does not preserve meaning. We will come back to some of these examples in discussing the syntax–semantics interface in formal semantics.

(1) a. John wants John to win ⟹ John wants to win
b. Everyone wants everyone to win ⟹ [!?] Everyone wants to win

(2) a. John voted for John ⟹ John voted for himself
b. Every candidate voted for every candidate ⟹ [!?] Every candidate voted for himself

(3) a. Three is even or three is odd ⟹ Three is even or odd
b. Every number is even or every number is odd ⟹ [!?] Every number is even or odd

The discovery that quantifiers, negation, and conjunctions were not properly interpretable at the level of Deep Structure as conceived in 1965 was one important factor in the genesis of the "Linguistic Wars". Generative Semanticists (Lakoff, McCawley, Ross, and others, cf., Huck and Goldsmith, 1995) responded by making Deep Structure "deeper", closer to "logical form", which for linguists then meant more closely resembling first-order predicate logic. Interpretive Semanticists (Jackendoff, Chomsky) kept the syntax closer to the "standard theory" and added additional interpretive mechanisms at various syntactic levels.

During all this time, the notion of semantic interpretation was in a rather primitive state. Katz, Fodor, and Postal worked with "semantic markers" modeled on phonological distinctive features, treating sentence meanings as bundles of semantic markers. The Generative Semanticists added first-order logic and uninterpreted but supposedly universal predicates and operators such as "CAUSE" and "BECOME". The reaction of philosophers of language was most notably formulated by David Lewis (1970).

> But we can know the Markerese translation of an English sentence without knowing the first thing about the meaning of the English sentence: namely, the conditions under which it would be true. Semantics with no treatment of truth conditions is not semantics. (Lewis, 1970, p. 1)

To linguists, concern with *truth* looked puzzling. Linguists were trying to figure out mental representations that could underlie linguistic competence.

"Actual truth" was (correctly) considered irrelevant, and truth *conditions* were not understood or appreciated.

1.2.2 Semantics in logic and philosophy before 1970

In the meantime great progress was made in semantics in logic and philosophy of language. The greatest foundational figure for formal semantics is Gottlob Frege (1848–1925). He made crucial advances to the analysis of variables and quantifiers, and introduced the distinction between *Sinn* and *Bedeutung*, sense and reference, the precursor of the modern intension–extension distinction. Frege is also credited with the Principle of Compositionality,[1] a cornerstone of formal semantics, and a principle quite universally followed in the design of the formal languages of logic, with a few interesting exceptions.[2]

> The Principle of Compositionality: the meaning of a complex expression is a function of the meanings of its parts and of the way they are syntactically combined.

Further advances were made by Russell, Carnap, and others. Wittgenstein (1922) first articulated the idea that "To know the meaning of a sentence is to know what is the case if it is true" (*Tractatus*, 4.024). Tarski (1902–1983) formalized *model theory* inside set theory and provided the first formal model-theoretic semantics for logical languages (Tarski, 1944; Feferman and Feferman, 2004); his goals concerned metalogic and the avoidance of semantic paradoxes, not semantics itself (Etchemendy, 1988).

Around that time, there was a major dispute within philosophy, the "Ordinary Language" vs. "Formal Language" war, concerning the role of natural language in philosophical argumentation, including the question of whether the analysis of natural language could be a source of philosophical insights, or whether natural language was too unruly and needed to be replaced by a suitably constructed formal language for purposes of exact argumentation. The Ordinary Language philosophers, including late Wittgenstein, Austin, Ryle, and Strawson, were a new generation who rejected the formal approach and urged closer attention to the functions of ordinary language and its uses, including much more attention to language in context, i.e., to pragmatics. In his critique of Russell in "On Referring", Strawson (1950) said, "Neither Aristotelian nor Russellian rules give the exact logic of any expression of ordinary language; for ordinary language has no exact logic." Russell rejected Strawson's critique, but added "I agree, however, with Mr. Strawson's statement that ordinary language has no logic" (1957). It is noteworthy that both sides in this "war" (as well as

[1] Not without some controversy; see Janssen (1983). And see Hodges (2001) for a discussion of the relation between compositionality and contextuality in Frege, and Pelletier (2001) for a third evaluation.

[2] It has been observed in a number of works (I learned it from Tarski, p.c., 1971) that the usual semantics for the quantifiers of first-order logic in terms of satisfaction and assignments is not strictly compositional.

Chomsky) were in agreement that logical methods of formal language analysis did not apply to natural languages.

The response of some formally oriented philosophers was to try to analyze ordinary language better, including its context-dependent features. Within philosophical logic, the foundational work of Frege, Carnap, and Tarski had led to a flowering of work on modal logic and on tense logic, on conditionals, on referential opacity, and on other philosophically interesting natural language phenomena. Quine had rejected modal and intensional notions as incurably unclear, but Kripke (1959), Kanger (1957a,b), and Hintikka (1962) revolutionized the field by providing a model-theoretic possible-worlds semantics for modal logic. And first Reichenbach (1947) and then Prior (1967) made great progress on the development of the logic of tenses, a notorious source of context-dependence in natural languages; Thomason (1996) identifies Prior as an important contributor to "natural language semantics logicism".

Paul Grice (1913–1988) contributed in a quite different way to the eventual resolution of the war. His work on conversational implicatures (Grice, 1975) showed that explanatory pragmatic principles can allow semantics to be simpler, so that the apparent gap between the "logicians' meaning," even in terms of standard, extensional first-order logic, of words like *the* and *if* and *or* and their ordinary-language meaning might be much less than had been supposed. And although not all of his proposals for the semantics of such words have survived, his methodological lesson and his pragmatic principles became highly influential for formal and informal semantics alike. Once his lessons sunk in, it became obligatory in studying semantics to also think about pragmatics, although whether and how pragmatics "belongs in the grammar" has remained controversial for decades (Horn, 2006).

1.2.3 Montague's work

Montague had been an important contributor to these developments in philosophical logic. Montague was a student of Tarski's, and at UCLA was a teacher and then a colleague of David Kaplan, co-authored a logic textbook with his colleague Donald Kalish, and was an active part of a strong logic group. As a logician, Montague built on the Frege–Tarski–Carnap tradition of model-theoretic semantics of logic and developed an intensional logic with a rich type theory and a possible-worlds model-theoretic semantics, incorporating certain aspects of "formal pragmatics", including the treatment of "indexical" words and morphemes like *I, you* and the present tense (Montague, 1968, 1970b). This was accomplished in part by treating both worlds and times as components of "indices", then generalizing from times to "points of reference", that is, complexes of relevant features of contexts, and treating intensions as functions from indices (not just possible worlds) to extensions. He generalized intensional notions such as property, proposition, individual concept, into a fully typed intensional logic, extending the work of Church (1951), Carnap (1956), and Kaplan (1964), putting together the function-argument structure common to type theories since Russell with the treatment of intensions as functions from indices to extensions.

In the late 1960s, Montague turned to the project of "universal grammar", which for him meant a theory of syntax and semantics encompassing both formal and natural languages, a groundbreaking project that became *Montague Grammar*,[3] and to which his last three papers were devoted, with plans for a book-length treatment interrupted by his untimely death in 1971.

That project evidently grew out of his work on the development of a higher-order typed intensional language suitable for doing philosophy. His paper "On the Nature of Certain Philosophical Entities" (NCPE) (Montague, 1969) contains a great deal that can be considered as much a matter of semantics as of philosophy, and foreshadows some of his work in his three final "language" papers. An important passage in that paper with respect to Montague's program occurs on pages 154–156, explaining his change from believing that philosophy should be done in the framework of set theory to believing that it should be done in the framework of intensional logic, and announcing his claim that he has constructed an adequate intensional logic.

> One system of intensional logic now exists which fully meets the objections of Quine and others, which possesses a simple structure as well as a close conformity to ordinary language, and concerning the adequacy of which I believe no serious doubts can be entertained. (Montague, 1969, p. 156)

This big "framework" change in Montague's approach to logic and philosophy is described and discussed in Cocchiarella (1981).

His attitude toward his work on natural language was ambivalent. On the one hand, he considered it worthwhile to demonstrate that "The syntax and semantics of certain not insignificant fragments of English can be treated just as formally and precisely as those of the first-order predicate calculus, and in very much the same manner" (Montague, in Staal, 1969, p. 274). However, at the same time, he asserted that "it would appear more important to extend the scope of constructed systems than to discover the exact rules of natural languages" (Staal, 1969, p. 275). As Thomason (1996) notes, Montague's quest for a "formal philosophy" grounded on his intensional logic remains unfulfilled and possibly quixotic, and his legacy is ironically rather in the "rather easy and not very important"[4] project of the analysis of ordinary language.

[3] The term entered the Oxford English Dictionary in 2002, the first citation being to Rodman (1972), a collection of papers by participants in a seminar taught at UCLA by the author.

[4] This unpublished quotation from Montague's notes, as well as evidence that Montague might have later revised his "rather easy" assessment, is discussed in Partee (2011, 2013).

Montague's work on the formal treatment of natural languages is all in his last three papers, "English as Formal Language" (EFL) (1970a), "Universal Grammar" (UG) (1970b), and "The Proper Treatment of Quantification in Ordinary English" (PTQ) (1973b). The one that had the most impact on linguists and on the subsequent development of formal semantics was PTQ: short, but densely packed (see Partee, 1975, 1997b; Janssen, 1994). Montague Grammar has often meant what Montague did in the fragment in PTQ and the extensions of PTQ by linguists and philosophers in the 1970s and 1980s. But it is the broader algebraic framework of UG that constitutes Montague's theory of grammar.

Before Montague, linguists took as the task of semantics the explication of ambiguity, semantic anomaly, and synonymy: the key questions were how many readings a sentence has, and which sentences share readings. The individuation of "readings" had always been problematic, though, and there was often crucial disagreement about data. Intuitions about "readings" undoubtedly rested in part on judgments concerning truth conditions, but truth conditions were never explicitly discussed. The methods used in early linguistic semantics primarily involved lexical decomposition in terms of *semantic features*, plus hypothesized abstract tree structures displaying scope relations and other aspects of *semantic structure*. The introduction of truth conditions as the basic semantic property of a sentence that a semantic theory should capture profoundly affected the adequacy criteria for semantics and led to a great expansion of semantic research.

Montague's (1970) paper, UG, contains the most general statement of Montague's formal framework for the description of language. The central idea is that anything that should count as a grammar should be able to be cast in the following form: the syntax is an algebra, the semantics is an algebra, and there is a homomorphism mapping elements of the syntactic algebra onto elements of the semantic algebra. In the PTQ grammar for a fragment of English, the syntax is not explicitly presented as an algebra, but if it were transformed into one, the elements would be the analysis trees.

The choice for the semantic elements is totally free, as long as they make up an algebra. The semantic elements, or *semantic values* as they are often called, could be taken to be the model-theoretic constructs of possible-worlds semantics as in Montague's fragments of English and most "classical" formal semantics, or the file change potentials of (Heim, 1982), or the game strategies of game-theoretical semantics, or the simple extensional domains of first-order logic, or hypothesized psychological concepts, or expressions in a "language of thought", or anything else; what is constrained is not the "substance" of the semantics but some properties of its structure and of its relation to syntactic structure. It is the homomorphism requirement, which is in effect the compositionality requirement, that provides the most important constraint on UG in Montague's sense, and it is therefore appropriate that compositionality is frequently at the heart of controversies concerning formal semantics.

Stokhof (2006) summarizes two important characteristics of Montague's theory as defined in UG and illustrated in PTQ:

a. Semantics is syntax-driven, syntax is semantically motivated. (Compositionality)
b. Semantics is model-theoretic.

A methodological principle implicit in Chomskyan syntax in the 1960s, encouraged although not required by the Katz–Postal hypothesis that meaning is determined at Deep Structure, and carried to extremes in Generative Semantics, was the principle that sameness of meaning should be reflected in sameness of syntactic Deep Structure. However, from the perspective of Montague Grammar, sameness of meaning does not require identity at any syntactic level (see Thomason, 1976): semantics is model-theoretic, not representational, not "syntactic".

In formal semantics, the core of sameness of meaning is sameness of truth conditions; and it does not require sameness on any syntactic level for two sentences to end up having the same truth conditions. Thus, formal semantics removed much of the motivation for that aspect of Generative Semantics. Its semantic goals could apparently be met, even exceeded, in a formal semantic approach, offering greater explicitness and compatible with more "conservative", less abstract, syntax. Thus, the rise of Montague Grammar was one factor in the decline of Generative Semantics and the fading away of the linguistic wars.

Details of Montague's analyses have in many cases been superseded, but in overall impact, PTQ was as profound for semantics as Chomsky's *Syntactic Structures* was for syntax. Emmon Bach (1989, p. 8) summed up their cumulative innovations thus: Chomsky's Thesis was that English can be described as a formal system; Montague's Thesis was that English can be described as an *interpreted* formal system.

1.2.4 Other contemporaneous work

While Montague clearly occupies a preeminent position in the history of formal semantics, he did not work in a vacuum. Works that influenced his thinking, as evidenced in his papers and in his seminars from 1967 and later, include, among others, Quine (1960b), Geach (1962, 1967) for puzzles of intensionality; Frege (1892b), Davidson (1965, 1967a, 1970), Kripke (1963), and various works of and/or conversations with Alfred Tarski, David Lewis, David Kaplan, Dana Scott, Rudolf Carnap, Alonzo Church, Yehoshua Bar-Hillel, J. F. Staal, Terence Parsons, and Barbara Partee, and several of his students including Hans Kamp, Dan Gallin, and Michael Bennett.

Independently, on the other side of the country, Donald Davidson and Gil Harman were both at Princeton from 1967 to 1969, interacting intensely, optimistic about the potential fruitfulness of linguistics–philosophy interactions and about the prospects of Generative Semantics, with its underlying structures somewhat resembling the structures of first-order logic. Together they produced some exciting conferences and the influential edited collection (Davidson and Harman, 1972). Davidson was also influential in urging a truth-conditional semantics for natural language, arguing from learnability that a finitely specifiable compositional semantics for natural languages must be possible (Davidson, 1967b). Davidson differed strongly from Montague in wanting to stay with first-order logic and to eschew such model-theoretic constructs as possible worlds and intensions. One of Davidson's contributions to semantics was the idea that the interpretation of most ordinary sentences includes an existentially quantified event argument, making sentences partly analogous to indefinite NPs (Davidson, 1967a); after the work of Terence Parsons (1990), the event argument became widely adopted.

There were major early contributions to formal semantics in Europe starting in the early 1970s. Renate Bartsch had come to UCLA to work with Montague just about at the time of his death; she and I had many fruitful discussions, but much more significant was her collaboration with Theo Vennemann, which began then at UCLA and continued in Germany (Bartsch, 1972; Bartsch and Vennemann, 1972b). Arnim von Stechow was an early and influential contributor to the rise of formal semantics in Germany and Europe, as well as to Generative Grammar and the integration of semantics and syntax (Kratzer et al., 1974; von Stechow, 1974). A number of formal semanticists[5] in other European countries point to von Stechow as the source of their earliest acquaintance with Montague's work. And influential contributions came from C. L. Hamblin in Australia (Hamblin, 1973) and M. J. Cresswell in New Zealand (Cresswell, 1973).

The earliest international conference on formal semantics (construed broadly) of natural language was organized by Ed Keenan at Cambridge University in 1973; eighteen of the twenty-five published contributions (Keenan, 1975) were by European scholars, including Östen Dahl, Hans Kamp, Peter Seuren, John Lyons, Renate Bartsch, Arnim von Stechow, Franz von Kutschera, Carl Heidrich, and Theo Vennemann. The biennial Amsterdam Colloquia, still a major forum for new results in formal semantics, started up in the mid-1970s and became international in the late 1970s. In the 1970s there were four textbooks on Montague grammar published in Germany, the last and best being Link (1979); all four were reviewed in Zimmermann (1981).

Other research that was not "Montague Grammar" in a narrow sense but which also fed into the development of formal semantics included Terence Parsons's combinatorial-based formal semantics of English (1972), David Lewis's influential paper (1970), which included a compositionally interpreted transformational fragment with a categorial grammar base, the work of the New Zealand philosopher and logician M. J. Cresswell (1973), who had

[5] Personal communication from several semanticists in Scandinavia and elsewhere; more details will be included in (Partee, in preparation).

been a participant in UCLA's famous "logic year" in 1967–1968, and the work of the American linguist Edward Keenan (1971a,b).

And there were of course other important early contributors to the development of formal semantics as well (for more on the early history of formal semantics, see Cocchiarella, 1981; Thomason, 1996; Partee, 1997b, 2011; Abbott, 1999; Cresswell, 2006; Stokhof, 2006; Janssen, 2011).

1.3 Central principles, issues, and points of divergence

1.3.1 The Principle of Compositionality and the theory-dependence of its key terms

At the heart of formal semantics are the principle that truth conditions form a core aspect of meaning and the methodologically central Principle of Compositionality.

For Montague, the compositionality requirement was the requirement of a homomorphism between syntax and semantics. The theory spelled out in Montague's UG (Montague, 1970b) requires that a grammar should take the following form: the syntax is an algebra, the semantics is an algebra, and there is a homomorphism mapping elements of the syntactic algebra onto elements of the semantic algebra. Montague's very general and formally very explicit definition leaves a great deal of freedom as to the nature of these algebras. For a logical language, the elements of the syntactic algebra can be the well-formed expressions. However, for a natural language, ambiguity makes that impossible, since the homomorphism requirement means that each element of the syntactic algebra must be mapped onto a unique element of the semantic algebra. So for a natural language, the elements of the syntactic algebra are usually taken to be expressions together with disambiguating structural descriptions, typically trees of some sort.

Differences among approaches can often be traced to three crucial theory-dependent terms in the Principle of Compositionality: "The *meaning* of a whole is a *function* of the meanings of its *parts* and the way they are syntactically combined." We discuss *meanings* in Section 1.3.2, the means by which meanings compose in Section 1.3.3, and *parts*, i.e., *syntax*, in Section 1.3.4, with more on the syntax–semantics interface in Section 1.4. The concept of *function* common to almost all semantic theories is the familiar set-theoretic one, but theories may differ in what further constraints they add on allowable functions.

1.3.2 What are meanings?

David Lewis provided a famous strategy for thinking about what meanings are: "In order to say what a meaning is, we may first ask what a meaning does, and then find something that does that" (1970, p. 22). There are different proposals about what to count as meanings (or as the linguistically relevant aspects of meanings) within formal semantics.

An *extensional language* is one in which compositionality holds for extensions: the extension of the whole is a function of the extensions of the parts and the way they are syntactically combined. In a simple case for such a language, the extension of a sentence is a truth value, the extension of a predicate expression (common noun phrase, simple adjective, verb phrase, predicative prepositional phrase) is a set of individuals; the extension of a proper name or a pronoun can be an individual. (See any formal semantics textbook, such as Gamut, 1991; Heim and Kratzer, 1998; and Chierchia and McConnell-Ginet, 2000).

Natural languages are at least intensional: the extension of the whole in some cases depends on intensions as well as extensions of parts, and the intension of the whole depends on the intensions of the parts.[6]

Within a decade, the importance of context-dependent expressions was increasingly recognized and led to increasing integration of formerly "pragmatic" matters into formal semantics (or "formal semantics and pragmatics"). Kaplan (1979) introduced the *character* of an expression, a function from contexts to intensions, and gave an account on which a sentence like *I am here now* is always true when spoken but is not a necessary truth. Kamp (1981b) and Heim (1982) formalized Karttunen's notion of a *discourse referent* introduced into the local context by the use of an indefinite NP (Karttunen, 1976) (see Brasoveanu and Farkas, Chapter 8) and built on Stalnaker's notion of a *common ground* established and constantly updated during the course of a conversation, resulting in two different but similar theories often collectively referred to as *dynamic* semantics,[7] treating meaning as a function from contexts to contexts. (See Section 1.5 below and Schlenker, Chapter 22.) An alternative conception of dynamic semantics was developed by Groenendijk and Stokhof (1990, 1991) and Veltman (1996), building in part on work by van Benthem, and with further linguistic development by Chierchia (1995a), and others.

[6] Because any two sentences with the same truth conditions have the same intension on standard possible-worlds analyses, intensions have often been said to be too weak to serve as meanings. There are various proposals for treating some constructions as *hyperintensional* and making use of a richer concept of *structured meanings* to handle them, an idea that has roots in Carnap's *intensional isomorphism* (Carnap, 1956) and in (Lewis, 1970) and was further developed in the pioneering work of Cresswell and of von Stechow (Cresswell, 1975, 1985; Cresswell and von Stechow, 1982; von Stechow, 1982; see also Duží et al., 2010). Montague formalized *intensions* of sentences as functions from possible worlds and variable assignments to truth values. And more generally, the intensions of all other categories of well-formed expressions are formalized as functions from possible worlds and variable assignments to the corresponding extensions.

[7] The common description of the Kamp–Heim approach as involving a dynamic conception of meaning may be challenged (M. Stokhof, p.c.) as a simplification or as misleading. In DRT, at least, it is the "construction rules" that build discourse representations that are dynamic; the model-theoretic interpretation, which involves the embedding of a discourse representation into a model, is static and classical. In the Amsterdam dynamic semantics of Groenendijk, Stokhof, Veltman, et al., classically constructed syntactic descriptions are dynamically interpreted.

In an independent refinement, Barwise and Perry (1983) and Kratzer (1989) argued for replacing possible worlds by (possible) *situations*, which for Kratzer are *parts* of possible worlds, enabling more fine-grained analysis of meanings (see Kratzer, 2011a).

There are other formalizations of what meanings are in contemporary formal semantics; see, for example, Hintikka's game-theoretic approach (Hintikka and Sandu, 1997), "Glue Semantics" (Asudeh and Crouch, 2002). Another important approach is that of of constructive type theory, as developed by Per Martin-Löf and applied to natural language semantics by Aarne Ranta (Ranta, 1995)[8]. See also Pagin, Chapter 3 in this respect.

1.3.3 How do meanings compose?

How are meanings put together? How does the compositional mapping from syntax to semantics work? The question of what sorts of functions are used to put meanings of parts together is inextricably linked to the questions of what meanings are, and of what count as syntactic parts. Frege (1892b) took the basic semantic combining operation to be *function-argument application*: some meanings are construed as functions that apply to other meanings. With a syntax such as categorial grammar providing the relevant part-whole structure, especially in the case of a simple extensional language, Frege's function-argument principle could be enough; with other kinds of syntax, other operations may be needed as well. (See Sailer, Chapter 21.) In the fragment in Montague's PTQ, the simple grammatical relations are all treated as intensional and are all interpreted by function-argument application (the function applies to the intension of the argument to give the extension of the result); some other special syntactic rules, especially ones that are partly "transformation-like," have more complex corresponding interpretation rules. Other possible approaches to meaning composition include unification (Bouma, 1988; Carpenter, 1992) and Glue Semantics, already mentioned. The possibility of type-shifting processes and various kinds of coercion complicates the picture and complicates the definition of compositionality (Pagin and Westerståhl, 2010a,b; Partee, 2007).

Formal semanticists also differ on whether a level of *semantic representation* is hypothesized to mediate between syntax and model-theoretic interpretation. Montague's own work exemplified both *direct model-theoretic interpretation* (1970b) and two-stage interpretation via translation into a language of Intensional Logic (1973b). Many but not all formal semanticists make use of an intermediate semantic representation in some formal language, either as a convenience ("for perspicuity," as Montague put it) or with the hypothesis that it might represent a "linguistically real" level of representation. Either approach can be compositional: a two-stage interpretation procedure is compositional if the syntax-to-semantic-representation mapping

[8] I am grateful to Martin Stokhof for bringing Ranta's work to my attention.

rules are compositional and the model-theoretic semantics of the representation language is also compositional. When those conditions are met, the intermediate language is, from a model-theoretic perspective, in principle eliminable. Linguists who hypothesize that it has some psychological reality would still want to keep such a level: it may represent an aspect of the means by which humans compute the mapping between sentences and their meanings. But it is a major challenge to find empirical evidence for or against such a hypothesis.

It is worth noting that it is possible to advocate direct model-theoretic interpretation without being anti-psychologistic, via the notion of *mental models* (Johnson-Laird, 1983). But approaches using mentally represented "formulas" (logical forms, conceptual representations) and computations on such formulas, as advocated by Jackendoff for many years, are preferred by many Chomskyan linguists.

Kamp in his Discourse Representation Theory (DRT) (1981b) proposed Discourse Representation Structures (DRSs) as a *non-eliminable* intermediate level of representation, with claimed psychological reality: Kamp hypothesized that his DRSs could be a common medium playing a role both in language interpretation and as objects of propositional attitudes. Kamp argued against full compositionality; he was challenged by Groenendijk and Stokhof (1991), who argued that a fully compositional *dynamic semantics* could accomplish what Kamp could do with DRT. Muskens (1993) proposed a reconciliation with his Compositional Discourse Representation Theory.

1.3.4 Part-whole structure: syntax

Logicians typically specify the syntax of a formal language as a recursive definition; in that case the requirement of compositionality as homomorphism can be satisfied by giving the semantics in the form of a parallel recursive definition. The "derivation trees" that correspond to the steps in applying a recursive definition become the elements of the syntactic and semantic algebras, and the homomorphism requirement says that each syntactic derivation must be mapped onto a unique semantic derivation.

The simplest linguistic examples take the elements of the syntactic algebra to be (sub)trees generated by a context-free grammar, and semantic interpretation to specify how the interpretation of a given subtree is computed from the interpretations of its immediate "daughter" subtrees; in this case the algebraic conception can be taken as compatible with and an improvement on what Fodor and Katz, and Katz and Postal, were aiming at with their Projection Rules.

Some formal semanticists argue in favor of having a *monostratal grammar* for the syntactic component, that is, a grammar with a single level of syntax and no transformations, such as Generalized Phrase Structure Grammar (GPSG), Head-Driven Phrase Structure Grammar (HPSG), Lexical-Functional Grammar (LFG), Categorial Grammar, Tree-Adjoining Grammar (TAG).

Others who work with a Chomskyan syntax take the syntactic input to the semantic component to be a specified level of semantically relevant syntactic representation called LF (a term meant to suggest "logical form", but defined purely theory-internally).

The relation between the preceding issues and syntax shows up clearly in debates about *direct compositionality*: some linguists argue that a directly compositional model-theoretic semantics can apply to non-abstract surface structures (see the debates in Barker and Jacobson, 2007), without abstract syntactic representations, movement rules, or a level of Logical Form. Advocates of direct compositionality use an enriched arsenal of semantic combining rules, including not only function-argument application but also function composition and a number of type-shifting operators. There may or may not be an inevitable tradeoff between optimizing syntax and optimizing semantics; it is a sign of progress that many linguists work on syntax and semantics with equal concern for both.

1.4 From Montague Grammar to contemporary formal semantics

1.4.1 Issues in combining formal semantics with syntax

The earliest works in Montague Grammar followed Montague's pattern of giving the syntax as a simultaneous recursive definition of expressions of all syntactic categories, and the semantics as a corresponding recursive definition of the interpretations of those expressions. In the first papers by linguists and philosophers, collected in Rodman (1972),[9] the form of the syntax was close to Montague's; Partee took on the challenge of trying to combine Montague Grammar (MG) with Transformational Grammar (TG) (Partee, 1973b, 1975), trying to recast transformational rules into recursive definitions that could be interpreted compositionally. On Partee's approach to the MG–TG combination, a syntactic derivation worked bottom-up with phrase structure rules as the basic building rules and with transformations applying whenever their conditions are met; Partee (1980) suggested that a grammar thus organized might be able to meet a constraint that all generated expressions be well formed, unlike the common practice in then-current transformational grammar to generate tree structures that had to be subjected to obligatory transformational rules to produce well-formed expressions.

The task of putting MG and TG together was made difficult by the fact that within TG, the "building blocks", or "kernel sentences", were closed sentences, while in MG, the syntax includes "indexed pronouns" (Montague's $he_0, he_1, \ldots$), interpreted as variables. (See later uses of PRO, *pro*, and traces in

[9] That first collection resulted from Partee's winter–spring 1972 seminar on Montague Grammar at UCLA, and contains papers by Partee, Bennett, Bartsch, Rodman, Delacruz, and others.

generative grammar.) One of Montague's innovations was the use of lambda abstraction as the central device involved in the interpretation of variable-binding constructions, possible only if well-formed expressions can contain elements interpreted as variables. And one immediate obstacle to synthesis was the existence of various "deletion rules" in TG. In classical TG, (4a) was derived from something like (4b) by "Equi-NP Deletion".

(4) a. Mary was eager to win.
    b. [$_S$ Mary was eager for [$_S$ Mary to win]]

But given the principle of compositionality, and given the way MG works by building up the meanings of constituents from the meanings of their subconstituents, this derivation seemed to present a problem. The syntactic derivation uses deletion, but the semantic derivation cannot: there is no permissible operation that would "delete" a piece of a meaning of an already composed subpart. Recall the discussion in Section 1.2.1 of the consequences of analyses like (4b) for sentences like (5a). The presumed Deep Structure (5b) would clearly give the wrong meaning.

(5) a. Everyone was eager to win.
    b. [$_S$ everyone was eager for [$_S$ everyone Tns win]]

The MG–TG resolution suggested in Partee (1973b, 1975) was that the "underlying" subject in the embedded sentence should be a bindable variable.[10] Partee followed Montague's line and bound it by lambda abstraction to make a VP type, as in (6a), assuming that the complement of the adjective *eager* is a VP. Others have proposed an S type for the infinitive, with the variable bound by the lambda abstract associated with the higher quantifier, as in (6b). In this very simple example, the VP in (6a) could alternatively just be base-generated and interpreted directly; Partee's "Derived VP rule" was motivated by VPs like *to see herself* or *to be elected*, which she derived transformationally from open sentences like she_0 *sees* her_0*self* and he_1 *elects* her_0.

(6) a. $[\![\text{to win}]\!] = \lambda x[\text{win}(x)]$
    b. alternatively: everyone′ (λx[x was eager for [x to win]])

In Chomskyan syntax, a corresponding change was eventually made. The first step, influenced by Reinhart (1976, 1983), involved replacing the "identical NP" by the special null element PRO, interpreted as a bound variable. A considerably later step, probably influenced by Heim (1982), introduced functional heads that could be interpreted as abstraction operators rather than assuming that indexed quantifiers themselves were responsible for binding. Other syntactic theories, like GPSG, HPSG, and LFG, and modern versions of Categorial Grammar, were developed after the quantifier and binding issues had become well known, and their design included mechanisms to deal with those problems.

[10] A similar proposal had already been made within Generative Semantics by McCawley (1968b).

There were other obstacles to combining Transformational Grammar and Montague Grammar, since some transformational rules were neither meaning-preserving nor easily reformulated as rules interpretable by a uniformly meaning-changing function, but the problem just described is a good example of an important problem in principle, whose resolution requires rethinking parts of the architecture of one theory or the other. In any case, the goal of combining TG with MG lost some of its urgency as various linguists began to realize that with a powerful semantics to do some of the work, some kinds of transformations – possibly all – might be eliminated.

Dowty (1978) argued in favor of eliminating all "governed transformations" (those whose application depended on the presence of a particular lexical class, such as the *easy*-class of adjectives implicated in the rule of "*Tough*-movement" mapping [structures corresponding to] *It's tough to solve this problem* to *This problem is tough to solve*) and replacing them by lexical rules which in effect transform the argument structure of a given lexical item (with or without some morphological change) and spell out the corresponding change in semantic interpretation.

Gazdar (1982) proposed that all transformations could be eliminated, so that the syntax could be context-free. He and colleagues developed his approach into Generalized Phrase Structure Grammar (GPSG) (Gazdar et al., 1985). Transformations were replaced by meta-rules that extended the grammar; for instance, in place of the passive transformation, a meta-rule specified how for each phrase structure rule generating a phrase containing a transitive verb and an object, plus possibly more constituents, a phrase structure rule should be added generating a corresponding passive verb phrase. And for each such meta-rule, there was a uniform semantic meta-rule mapping the compositional interpretation rules for the original phrase structure rules onto compositional interpretation rules for the derived phrase structure rule. So instead of a transformation mapping active sentences to passive ones, there would be a meta-rule mapping certain VP rules onto other VP rules, with corresponding semantics. This was the first of several proposals for a non-transformational syntax with a compositional formal semantics. Extended Categorial Grammar (Bach, 1981, 1984) was developed around the same time, and HPSG followed soon after (Pollard and Sag, 1994). LFG had been invented a little earlier (Kaplan and Bresnan, 1982); a compositional formal semantics interpretive component for it was proposed by Halvorsen (1983), and in the intervening years there have been other proposals such as the earlier mentioned Glue Semantics.

1.4.2 Case study: noun phrases, quantifiers, and relative clauses

Much of the discussion in this article has been at an abstract level; in this section we will look at some examples that illustrate a number of issues.

One of the reasons that so many philosophers and linguists had agreed earlier that linguistic structure and logical structure were very different was the apparent mismatch between the syntax and semantics of noun phrases (NPs).[11] Russell considered it an illogicality of English that expressions semantically so different as *Jones*, *a philosopher*, *every student*, *no man*, and *the king* have largely the same syntactic distribution,[12] and thus evidently belong to the same syntactic category (NP).

A major legacy of PTQ was the very important and influential analysis of noun phrases as uniformly denoting generalized quantifiers. Part of the appeal of this analysis for linguists was that it captured the important semantic differences among NPs headed by different determiners, as in generative semantics treatments, while at the same time giving all NPs a similar syntactic structure and an interpretation of the same semantic type, interpreting them all as generalized quantifiers, denoting sets of properties of entities (see Westerståhl, Chapter 7). Because most linguists had earlier known nothing about type theory, certainly nothing about generalized quantifiers, those who wanted to capture meaning at Deep Structure had been led to posit abstract deep structures that resembled first-order logic; dependence on first-order logic had made it impossible for linguists to imagine giving an explicit semantic interpretation for *the* or *a/an* or *every* or *no* that didn't require decomposition into formulas with quantifiers and connectives, more or less like the translations one finds in logic textbooks. The Generative Semanticists embraced such structures and made underlying structure look more like first-order logic (Lakoff, 1971a,b; McCawley, 1970), while Chomsky and Jackendoff rejected putting such logical decompositions into the syntax and devised various proposals for how some sort of semantic component could interpret the combination of deep and surface structure (Chomsky, 1971; Jackendoff, 1972). One can speculate that the rift might never have grown so large if linguists had known about generalized quantifiers earlier; the productive teamwork of Barwise and Cooper (1981) is a classic early example of how formal properties and linguistic constraints and explanations can be fruitfully explored in tandem with the combined insights and methodologies of model theory and linguistics, and generalized quantifiers have continued to be a fertile domain for further linguistically insightful work exploiting formal tools (Peters and Westerståhl, 2006; Szabolcsi, 2010; Keenan and Paperno, 2012).

In recent decades the differences among different noun phrase interpretations, including differences among referential, predicative, and quantificational uses, have led many semanticists to a less uniform treatment of the semantics of NPs, to the exploration of type-shifting mechanisms to help keep the best properties of a uniform analysis while doing justice to the apparent flexibility of natural language interpretation (Partee, 1986; Hendriks, 1993; Szabolcsi, 1997), and to innovative proposals for many aspects

[11] I use the older term NP in a broad sense, to include the contemporary syntactic categories NP and DP.

[12] If I say *Scott was a man*, that is a statement of the form "x was a man", and it has *Scott* for its subject. However, if I say *the author of Waverley was a man*, that is not a statement of the form "x was a man" and does not have *the author of Waverley* for its subject. (Russell, 1905, p. 488)

of NP (or DP) interpretation. Recent and current work focuses on such topics as the mass-count distinction, plurality, weak and strong quantification, indefinites, vagueness and gradability among modifiers and quantifiers, modal and temporal aspects of NP and DP interpretation, and more. (See Dekker and Zimmermann, Chapter 6; Westerståhl, Chapter 7; Brasoveanu and Farkas, Chapter 8; Nouwen, Chapter 9; and Cohen, Chapter 10.) Quantification and related issues in the interpretation of NPs and DPs have also proven to be extremely fertile ground for cross-linguistic and typological studies in formal semantics (Bach et al., 1995; Chierchia, 1998b; Kratzer, 2005; Matthewson, 2001; von Fintel and Matthewson, 2008).

Here we give an illustration of the methods of formal semantics and their impact on the resolution of the "mismatch" between logical and linguistic structure by considering one aspect of the analysis of restrictive relative clauses like *that Pat had lost* in (7a) and (7b) and their interaction with the semantics of various quantifiers and determiners.

(7) a. Mary found a hat that Pat had lost.
    b. Mary found every hat that Pat had lost.

In the 1960s, there were debates about whether the relative clause combines with the common noun phrase (today's NP), as in structure (8), or with the full DP *a hat, every hat*, as in structure (9).

(8) Mary found $[_{DP}$ a/every $[_{NP}$ $[_{NP}$ hat][that [Pat had lost ___]]]]

(9) Mary found $[_{DP}$ $[_{DP}$ a/every $[_{NP}$ hat]][that [Pat had lost ___]]]

There were also debates about the semantics of the relative clause, with some arguing that in (7a) *that Pat had lost* means *and Pat had lost it*, whereas in (7b) it means *if Pat had lost it*, creating tension between the uniform surface structure of *that Pat had lost* in (7a) and (7b) and the very different "underlying" semantic interpretations posited for them (see Stockwell et al., 1973), inspired by the structure of their translations into first-order logic, as in (10a) and (10b).

(10) a. $\exists x(\text{hat}(x)\ \&\ \text{lost}(\text{Pat}, x)\ \&\ \text{found}(\text{Mary}, x))$
     b. $\forall x((\text{hat}(x)\ \&\ \text{lost}(\text{Pat}, x)) \rightarrow \text{found}(\text{Mary}, x))$

The formal semantics perspective suggests searching for a unitary syntax and meaning for *that Pat had lost* and locating the semantic difference between (7a) and (7b) in the semantics of *a* and *every*. The solution (due to Quine (1960b) and Montague (1973b)) requires structure (8): the noun and relative clause denote sets, and their combination denotes the intersection of those two sets. Then the phrase *hat that Pat had lost* denotes the set (11a), whose characteristic function is denoted in the lambda-calculus by (11b).

(11) a. $\{x : x \text{ is a hat and Pat had lost } x\}$
     b. $\lambda x.\text{hat}(x)\ \&\ \text{lost}(\text{Pat}, x)$

Different theories of the semantics of determiners give different technical implementations of the rest of the solution, but that first step settles both the syntactic question and the core of the semantics. Nouns and common noun phrases (NPs) denote sets (or their characteristic functions), and restrictive relative clauses also denote sets (or their characteristic functions); in extensional type theory, they are both of type $\langle e, t\rangle$. Restrictive relative clauses uniformly combine with an NP to give a new NP; the semantics is just set intersection (12) or its equivalent in the lambda calculus as shown above in (11b).

(12) $\|[_{\text{NP}}\ \text{NP REL}]\| = \|\text{NP}\| \cap \|\text{REL}\|$

In the treatment of DPs as generalized quantifiers, the determiner or quantifier is interpreted as a function that applies to a set and gives as a result a generalized quantifier of type $\langle\langle e, t\rangle, t\rangle$, a set of sets of individuals.[13] In the classic treatments of Montague and of Barwise and Cooper, the interpretation of a DP of the form *a NP* is the set of all those sets that have a non-empty intersection with the *NP* set, and the interpretation of a DP of the form *every NP* is the set of all subsets of the set denoted by *NP*.

As a result, sentence (7a) asserts that the set of hats that Pat had lost and the set of hats that Mary found overlap; (7b) says that the set of hats that Pat had lost is a subset of the set of hats that Mary found. Thus the apparent difference in the interpretation of the relative clause in the two DPs turns out to be the predictable result of the semantic difference between the two determiners; there is no need to give the relative clause a non-uniform interpretation, and no reason to give the DPs syntactic structures resembling the formulas of first-order logic. See Partee (1995) for a fuller argument, and Barwise and Cooper (1981) for a fuller treatment of the classical formal semantics of determiners.

1.4.3 Varieties of syntax and semantics combinations

As noted above, formal semantics can in principle be combined with many kinds of syntax, but different syntactic frameworks may require differences in how the semantics is structured if compositionality is to be observed. We illustrate with a thumbnail review of how quantifier scope ambiguities have been handled in a number of different approaches, omitting most details.

Quantifier scope ambiguity is a major challenge for compositionality. A sentence like (13) is semantically ambiguous, but there has never been much evidence for it being syntactically ambiguous; the challenge is to try to find an analysis which is both compositional and syntactically well motivated.

[13] Within generalized quantifier theory, this functional treatment of determiners can be equivalently replaced by a treatment of determiners as denoting a relation between two sets; the two approaches are interdefinable. The relational interpretation is often logically more perspicuous, the functional treatment more faithful to natural language compositional structure.

(13) Every student read one book.

There were at least six kinds of solutions proposed from the 1960s to the 1980s, even before the introduction of choice functions and various non-quantificational analyses of indefinites, and in the ensuing decades more proposals have been made than I can list.[14]

Generative Semantics

The first serious attempts to account for quantifier scope ambiguity in generative grammar came from Generative Semantics; there was great progress in uncovering some of the principles that govern possible quantifier scope, bound variable anaphora, and related phenomena. Classic works include Bach (1968); Lakoff (1971b); McCawley (1970). Generative Semanticists proposed underlying structures that looked similar to the structure of first-order logic, plus a transformation of Quantifier Lowering so that a quantifier that starts as a sentence operator can end up as a determiner on an NP. The actual data about scope possibilities were controversial, as pointed out by Carden (1976), who was an early advocate of and pioneer in experimental methods in semantics. The perceived need to constrain derivations so that scope in English corresponds to surface c-command (now considered incorrect for English but probably correct for some languages) led to transderivational constraints (Lakoff, 1973), observed by Langendoen (2001) to be an early example of optimality-theoretic-like devices.

Interpretive Semantics

Interpretive Semanticists, led by Jackendoff and Chomsky, maintained an "autonomous syntax" and argued that different semantic phenomena were to be accounted for at various different syntactic levels: argument structure is determined at Deep Structure, but quantifier scope and variable binding may depend on structures at various levels, possibly including surface structure. Classic works include Chomsky (1971); Jackendoff (1972). Cooper and Parsons (1976) showed how a basic version of the scope mechanisms of Generative Semantics, Interpretive Semantics, and Montague Grammar (see below) were intertranslatable.

Montague's Quantifying-In rule

Montague's rule-by-rule approach is illustrated by the rule of "Quantifying-In" in Montague (1973b), which combines an NP with a sentence *S* by

[14] Martin Stokhof asks how this proliferation of proposals sits with the assumed empirical nature of the problem. Indeed, the quantifier scope problem is puzzling precisely because there is no independent debate about the syntax; in this way it differs from constructions for which the demands of compositionality can help to constrain choices among syntactic analyses, as with the attachment of relative clauses or the internal structure of comparative constructions. The choices in this case are between different kinds of theoretical apparatus in the syntax and semantics, and evaluation is of necessity theory-wide more than construction-specific.

substituting that NP for an indexed pronoun (interpreted as a variable) in *S* and substituting appropriate pronouns for further occurrences of the same indexed pronouns. The given NP is semantically interpreted as a generalized quantifier, and that generalized quantifier takes as argument the property expressed by a representation obtained by lambda-abstracting over the corresponding variable in the formula corresponding to *S*. For sentence (13), there are different derivations that differ in the order in which the two NPs are quantified into the structure, with correspondingly different scopes in the semantics. On this approach, a single surface syntactic structure may be derived via distinct derivations; compositionality is homomorphism between syntactic and semantic derivations, not between some levels of syntactic and semantic representations. But as Cooper and Parsons (1976) showed, one could algorithmically convert a Montague-style analysis into a Generative Semantics or Interpretive Semantics analysis, and vice versa (in this particular domain).

Cooper storage

Robin Cooper (1975) proposed an approach that would treat sentence (13) as syntactically unambiguous while analyzing it as semantically ambiguous, and maintaining a quasi-compositional semantics. His idea was to amend the syntax–semantics interface so as to non-deterministically derive a *set* of interpretations for each syntactic structure; interpretation was compositional apart from the element of non-determinism. In the process of semantic interpretation (bottom-up, like the syntactic derivation), when you hit a quantifier phrase, you optionally "store" it, and then you may "retrieve it" from storage when you hit a suitable higher node, such as an S or a VP, at which point you interpret the combination of the NP with the S or VP as in Montague's treatment. Scope islands represent points in the derivation where the storage register must be empty. It is of interest that the monostratal syntactic theory GPSG (Gazdar et al., 1985) uses a context-free grammar with a semantics that is quasi-compositional in just this way: straightforwardly compositional but with the use of Cooper storage to interpret sentences like (13).

Quantifier Raising

Later versions of Chomskyan generative grammar added a syntactic level of "LF" or "Logical Form", intended to provide the syntactic input to semantic interpretation. One early proponent of such a level was Robert May (May, 1977, 1985). His syntactic rule of Quantifier Raising, roughly inverse of the Generative Semanticists' Quantifier Lowering, produces a derived syntactic structure in which the various quantifiers are adjoined to the clause that represents their immediate scope. In this and a number of other respects, the LF approach may be approximately regarded as "Generative Semantics upside down".

Type-shifting

David Dowty, whose dissertation followed the Generative Semantics approach but who soon became a leading Montague Grammarian and formal semanticist, suggested in the 1970s that many transformations, especially "governed" ones, should be eliminated and replaced by lexical rules (Dowty, 1978, 1979). This is a very early example of the idea that much of the "grammar" is really contained in the lexicon, and rules that map lexical items onto related lexical items can often be interpreted semantically in terms of diathesis changes. Herman Hendriks (Hendriks, 1988) applied this perspective to quantifier scope ambiguity. His analysis does not require any "movement"; he derives alternative readings via type-shifting of verbs and other functors so as to change relative "function-argument command" relations, extending the kinds of "function-argument flip-flop" shifts introduced by Partee and Rooth (1983). Relative scope is then reflected in local function-argument hierarchical structure rather than requiring some level with attachment to higher S or VP nodes. Using a type-shifting approach to quantifier ambiguity rather than movement rules is one good example of direct compositionality, as discussed in Section 1.3.4.

Underspecification

For Montague, relative scope was captured at the level of derivation trees. Muskens (2004), influenced by work in computational linguistics on underspecification of the parsing process, and by Reyle's work (1993) on underspecification in Discourse Representation Theory, takes that idea a step further to provide a formalism that underspecifies syntactic and semantic ambiguity analogously, with the help of descriptions in the object language. Muskens provides underspecified derivation trees with constrained possibilities for how they may be completed. Each corresponding complete derivation tree generates the given sentence with one of its possible readings. One of the appeals of this approach, not so far taken up in the psycholinguistic literature as far as I know, is its potential for solving the problem of the "psychological reality" of quantifier scope ambiguity. Not only is there no perceived syntactic ambiguity in a sentence like (13), but there is little evidence of the kind of combinatorial explosion of ambiguity that is otherwise predicted for sentences with more than two quantifiers, and little evidence that ordinary speakers are very sensitive to the large numbers of quantifier scope readings that linguistic analyses have classically predicted. Underspecification might very well be a psychologically as well as computationally reasonable approach: the information is "there", but it need not always be "computed" in actual processing.

Pseudoscope

All of the approaches discussed above were designed to account for quantifier scope ambiguity in its "classical" form, with the distinct readings the same as in standard predicate logic. The differences all concerned how the syntax–semantics interface works to generate the sentences and derive the readings. But starting in the 1980s, with some even earlier suggestions, the nature of the readings was challenged, particularly but not only for indefinite NPs like *a book*. The relation between semantic wide scope and semantic or pragmatic notions of *specificity* was questioned by linguists in several frameworks (Fodor and Sag, 1982), and the idea that indefinite NPs are always to be interpreted via existential quantifiers was particularly strongly challenged in the work of Kamp and Heim discussed in Section 1.5.2 below. The importance of non-quantificational analyses of indefinites was emphasized in Kratzer (1998b); see also Brasoveanu and Farkas, Chapter 8.

1.5 The increasing role of context: formal semantics and pragmatics

1.5.1 Early views on the autonomy of context-independent semantics

The term *pragmatics* is due to the philosopher and semiotician Charles Morris (1938). Within semiotics, the general science of signs, Morris distinguished three branches: syntactics (now syntax), the study of "the formal relation of signs to one another"; semantics, the study of "the relations of signs to the objects to which the signs are applicable" (their designata); and pragmatics, the study of "the relations of signs to interpreters" (Morris, 1938, p. 6). Semantics in the work of Formal Language philosophers such as Russell is pure semantics, just as is the semantics of first-order logic; no consideration was paid to indexicals or demonstratives. The need for attention to context-dependence in interpretation was one of the concerns of the Ordinary Language philosophers.[15] One of the first aspects of context-dependence on which logicians made progress was tense, where Prior and Reichenbach were pioneers. And as noted in Section 1.2.3, Montague was an important contributor to these developments, developing a higher order typed intensional logic that unified tense logic and modal logic, and more generally unified "formal pragmatics" with intensional logic.

In early Montague Grammar (e.g. as reported in Partee, 1984a), context-dependence was recognized and treated by the use of free variables. Montague's own PTQ included indexed pronouns like he_3; these were normally replaced by ordinary gendered pronouns in the course of a derivation as part of some rule that involved variable binding, but they could also remain

[15] As Martin Stokhof has reminded me, the kind of context-dependence of interpretation that Ordinary Language philosophers were concerned with was much wider ranging than the type of context-dependence that is exhibited by indexicals. I neglect the deeper issues here.

free, in which case it was suggested that they could provide the interpretation for pronouns without antecedents (demonstratives, indexicals, and anaphoric pronouns without sentence-internal antecedents), to be interpreted by a variable assignment considered as part of a context of use. Other early examples of "free variables" to be interpreted via assignment functions coming from the context of use came from work on tense and aspect that appealed to one or sometimes two or more temporal indices corresponding to "now" and possibly to additional "reference times" (Kamp, 1971, 1979; Vlach, 1973). Implicit arguments were often treated as covert free variables, and a free relation variable was suggested as part of the interpretation of genitive constructions like *Mary's team* (Barker, 1995; Partee, 1997a).

In such work, the assumption was that the output of semantic interpretation is something like a Kaplanian "character," a function from contexts to an intension. This makes the semantics autonomous in the sense that the compositional interpretation process is context-independent, and a complete sentence can be interpreted semantically, and then "handed over" to a pragmatic component to fill in the values of the context-dependent variables. For tense and other indexical expressions, such an approach seems perfectly appropriate: we understand the invariant meaning of *I am alone here now*, and we know how to take contextual information to fill in the values of speaker, time, and place to arrive at a specific proposition.

But by the early 1980s, evidence of more complex interactions between semantic interpretation and context-dependence was accumulating to the point where Lewis's dictum, "In order to say what a meaning is, we may ask what a meaning does, and then find something that does that" (1970, p. 22), called for considering some aspects of context-dependence, traditionally thought of as pragmatics, as integral parts of meaning. What is generally referred to as *dynamic semantics* was a first result.

1.5.2 The beginnings of dynamic semantics

A landmark innovation in the integration of context-dependence into semantics came with the dynamic semantics[16] of Kamp (1981b) and Heim (1982). Consider the contrasting mini-discourses in (14) and (15):

(14) A baby was crying. It was hungry.

(15) Every baby was crying. #It was hungry. ("#" means "anomalous".)

On the Kamp–Heim theory, an indefinite NP like *a baby* in (14) introduces a "novel discourse referent" into the context, and the pronoun *it* in the second sentence of (14) can be indexed to that same discourse referent, whose "life span" can be a whole discourse, not only a single sentence. The discourse

[16] An even earlier example of interpretation via a context that is systematically built from denotations of expressions can be found in Groenendijk and Stokhof (1979).

referent introduced by an essentially quantificational NP like *every baby* in (15), however, cannot extend beyond its clause, so the pronoun *it* in (15) is anomalous. The Kamp–Heim theory also includes an account of the famous "donkey"-sentences of Geach, variants of which are given in (16a) and (16b).

(16) a. If a farmer owns a donkey, he always beats it.
b. Every farmer who owns a donkey beats it.

These sentences had previously resisted compositional analysis, including with the tools of Montague Grammar and its early extensions. On Kamp's and Heim's theories, the indefinite *a donkey* introduces a *discourse referent* into the local context, but has no quantificational force of its own; it ends up being bound by the "unselective quantifiers" *always* in (16a) and *every* in (16b). The theories involve the interdependent areas of quantification and anaphora, and the relation of sentence semantics to discourse semantics and pragmatics, giving rise to much new work in these areas. See Brasoveanu and Farkas, Chapter 8.

1.5.3 Further developments in formal pragmatics and semantics

In subsequent work in formal semantics, the treatment of context-dependent phenomena has been a central concern. Other aspects of formal pragmatics such as the treatment of presupposition and focus-sensitivity have also long been recognized as central to the study of meaning, so that "formal semantics and pragmatics" has increasingly come to be seen as a single field, as described in Kadmon (2001) and as evidenced by the online journal *Semantics and Pragmatics* founded in 2007. The integration of pragmatics and semantics has been further supported with arguments that the interpretation of pragmatic implicatures occurs interspersed with semantic interpretation (Chierchia, 2004).

Information structure has also played an important role in the development of formal pragmatics; notions such as topic, focus, presupposition, and discourse structure have been central concerns of linguists and philosophers of many different theoretical inclinations since before Montague's work, and have loomed large in discussions of pragmatics and semantics and the border between them (see Beaver, 1997; Karttunen and Peters, 1979; Roberts, 1998; Rooth, 1992; von Heusinger and Turner, 2006). The formalization of Gricean implicatures and Potts's (2005) proposal to treat conventional implicatures as part of semantics have also led to active research areas.

Semantics and pragmatics are both involved in recent work on Optimality-theoretic semantics and bidirectional OT, as well as in game theoretic approaches; these and other approaches to formal pragmatics are not well represented in this book, but see Hintikka and Sandu (1997), Chierchia (2004), Benz et al. (2005), Blutner et al. (2006), von Heusinger and Turner (2006). (See also Schlenker, Chapter 22.)

1.6 Formal semantics and psychology: the Frege–Chomsky conflict

The relation of linguistics to psychology is nowhere more controversial than in semantics. While it is easy enough to argue that the syntax of a person's language is directly a function of what is in that person's head, the parallel claim for semantics is highly controversial. The clash between Frege and Chomsky on the issue of "psychologism" and the notion of semantic competence has led to a conflict in the foundations of formal semantics, sometimes expressed as the question of whether semantics is a branch of mathematics or a branch of psychology (Partee, 1979, 1988b; Thomason, 1974).

Frege, whose ideas were part of the foundation of Tarski's and Montague's work, took an anti-psychologistic view of meanings, and so did many other logicians and philosophers, including Russell (sometimes), Tarski, Carnap, and Montague. Their tradition of "objective" (though abstract) meanings contrasts with the psychologistic view of meanings "in the head" (Fodor, Lakoff, Jackendoff, and all psychologists). The psychologistic view fits into the Chomskyan tradition of focusing on linguistic competence, defined in terms of the implicit knowledge of a native speaker.

Most formal semanticists who are linguists are very much concerned with human semantic competence. Some formal semanticists have advocated following David Lewis (1975b) in distinguishing semantics from knowledge of semantics, making semantic competence interestingly different from syntactic competence (Partee, 1988b).

In recent decades the question of whether meanings are "in the head" has taken an interesting turn as a result of work in the philosophy of mind. The early arguments by Kripke and Putnam (Kripke, 1972; Putnam, 1975) against a purely "internalist" account of meaning were arguments based on a narrow view of psychological states common at the time, a view later epitomized by Fodor's *methodological solipsism* (1980). However, philosophers working on the philosophy of mind in subsequent decades have increasingly argued that much of psychology, including perception as well as belief and other psychological states, needs to be viewed relationally, concerning the individual and her environment (Burge, 1992, 2003; Stalnaker, 1989). As Stalnaker puts it, meanings may be "in the head" in the way that footprints are "in the sand"; it's a perfectly reasonable way of speaking, but a footprint clearly isn't a footprint without a certain relation to causal history, and meanings probably aren't either. What is semantic competence? For formal semanticists, it is common to take the fundamental characterization of semantic competence to involve knowledge of truth conditions: given a sentence in a context, and given idealized omniscience about the facts concerning some possible situation, a competent speaker can judge whether the sentence is true or false in that situation. From that basic competence, allowing idealizations about computational capacity, it follows that a competent native speaker can also make judgments about entailment relations between sentences; so idealized semantic competence is widely considered to consist in knowledge of truth conditions and entailment relations of sentences of the language.

Many linguists have resisted the relevance of truth conditions and entailment relations to natural language semantics, and some still do. Some of the objections to truth conditions are countered by arguing that part of human semantic competence is the matching of sentences with their truth conditions relative to possible worlds (including fictional worlds), with no necessary reference to the "actual world".

Entailment is a central semantic concern in logic and remains so in formal semantics. Given the distinction between meaning and knowledge of meaning, there is no contradiction between taking entailment relations as central to meaning while acknowledging that human language users are not logically omniscient and hence do not always know whether a given sentence of their language entails another. In this sense it may not be unreasonable to say that speakers can't know their language; logical fallibility may be considered a sort of "performance limitation". The real problem arises if the semantic value of a proposition is taken to be its truth conditions, and propositions are analyzed as the objects of verbs like *believe*; then we get the unwelcome conclusion that sentences with the same truth conditions should be substitutable in belief-contexts. This is the problem of *hyperintensional contexts* (see Pagin, Chapter 3). And it is this problem that has motivated richer notions of meaning such as the structured meanings mentioned earlier.

Cognitive semanticists replace concern with logical entailment by concern with human inference; formal semanticists see the relation of entailment to actual human inference as indirect. However, semanticists of many stripes agree about the importance of revising the formal logics invented by logicians to model the "natural logic(s)" implicit in the semantics of natural languages.

A number of linguists still think of semantics in terms of a "level of representation" of expressions analogous to a syntactic or phonological level. A representational view of semantics is quite congenial to the popular computational theory of mind. But David Lewis's objection that semantics without truth conditions is not semantics remains in force, so that formal semanticists consider such representational views at best incomplete. The contrasting model-theoretic view that is dominant in formal semantics takes semantic interpretation to relate expressions to elements of models (possibly mental models; see Johnson-Laird, 1983) defined in terms of constituents such as possible situations, entities, properties, truth values. Whether some level of representation plays an intermediate role in interpretation, as in the two-stage method of interpretation discussed in Section 1.3.3, remains an open question that so far is not obviously an empirical one.

For further issues in the relation of formal semantics to cognitive science, see Baggio, Stenning and van Lambalgen, Chapter 24.

1.7 Other issues

1.7.1 Natural language metaphysics, typology and universals

The issue of "Natural Language Metaphysics" (Bach, 1986b) is an important foundational area that connects cognitive issues with formal semantics and linguistic typology. What presuppositions concerning the constitution and structure of the world as humans conceive it are built into human languages, and how, and which are universal? These questions may concern both semantic structure and semantic content, from the semantic difference between nouns and verbs to the content of color terms or whether time is discrete or continuous. Their investigation may challenge the lines between semantic knowledge and other kinds of knowledge. Formal semantics, following the logical tradition, initially employed relatively "austere" model structures; recent investigations, particularly into lexical semantics, have led to richer model structures.

Formal semantics has steadily expanded its scope to include work on a wide variety of languages, and semantic fieldwork is an increasingly active and well-developed field. Semanticists have made significant contributions in recent decades to the study of linguistic universals and linguistic typology (Bach et al., 1995; Bittner, 1994; Keenan and Paperno, 2012; Matthewson, 2010; von Fintel and Matthewson, 2008). Typological investigations, especially but not only in the domains of quantification and of tense, aspect and mood, have included the distribution of nominal vs. adverbial quantification, the semantics of languages that lack a determiner category, and the range of tense and aspect systems. Recent work on the formal semantics of sign languages has led to enriched perspectives on the semantics of demonstratives, variable binding, propositional attitudes, quotation, and other domains (Quer, 2005; Schlenker, 2011e).

1.7.2 Experimental, corpus-based, and computational formal semantics and pragmatics

Respected early work on applications of formal semantics to computational linguistic tasks included some in North America (Gawron et al., 1982; Halvorsen, 1983; Hobbs and Rosenschein, 1978; Schubert and Pelletier, 1982) and some in Europe (Cooper, 1987; Hirst, 1986; Scha, 1976). One of the most successful and large-scale efforts was the Rosetta project led by Jan Landsbergen at Philips Eindhoven in the 1980s (Rosetta, 1994), which grew out of earlier work by Landsbergen and by Scha. Recently, there have been advances in methods for making formal semantics computationally tractable, and interest in computational formal semantics has increased to the point that there are textbooks and courses in the field (Blackburn and Bos, 2003, 2005; Bunt and Muskens, 1999). A few companies have begun to offer computational linguistic products that use formal semantics and others are exploring the possibility of doing so. For some kinds of natural language processing problems, such as the solving of the kinds of logic puzzles found on the Graduate Record Exam (GRE), formal methods can offer important advantages and have been implemented (Lev et al., 2004). Connecting computational linguistics, psychology, and formal semantics, there is current research aimed at combining the best of statistical/probabilistic and formal approaches, exemplified by the work of Joshua Tenenbaum's Computational Cognitive Science group at the Massachusetts Institute of Technology (Tenenbaum et al., 2011).

There has always been an interest in acquisition and processing of semantics, often together with syntax, and in recent decades there has been a significant amount of research in these areas specifically connected with formal semantics and pragmatics (see, for instance, Chambers et al., 2002; Chierchia et al., 2004; Crain et al., 1996; Crain and Thornton, 2011; Frazier, 1999; Papafragou and Musolino, 2003; Philip, 1995). The emerging area of "experimental pragmatics" is represented by (Noveck and Sperber, 2007; Sauerland and Yatsuhiro, 2009).

Computational and experimental methods in and applications of formal semantics and pragmatics are developing very quickly, and neither this introductory chapter nor this handbook is pretending to try to do justice to them.

Acknowledgments

I am grateful to many colleagues for suggestions and responses to queries along the way, including Reinhard Muskens, Richmond Thomason, Nino Cocchiarella, Robert Stalnaker, and Seth Yalcin. I am particularly grateful for extensive helpful comments from a referee who agreed to be non-anonymous so that I could thank him publicly, Martin Stokhof.

2

Lexical semantics

James Pustejovsky

2.1 Introduction

Lexical semantics is the study of what words mean and how their meanings contribute to the compositional interpretation of natural language utterances. The lexicon can be seen as that component of the grammar that encodes both the information required for composition in the syntax and the knowledge for multiple levels and types of semantic interpretation. Lexical entries are richly structured objects that act as triggers both to compositional operations and to entailments and implicatures in the context of larger discourses. Because any semantic interpretation requires access to knowledge about words, the lexicon of a grammar must provide a systematic and efficient way of encoding the information associated with words in a language.

Four key questions arise when determining how to model the meanings conveyed by words:

(i) What are the semantic components that constitute word meaning?
(ii) How are word meanings differentiated and how are they related to each other?

(iii) How does the meaning of individual words drive the compositional process to make semantically coherent sentences?
(iv) When is a component of word meaning considered "lexical" rather than "world" knowledge?

As the linguistic phenomena associated with lexical semantics become better understood, several theoretical assumptions have emerged across most current models of word meaning. These can be summarized roughly as follows:

- Lexical meaning involves a kind of componential analysis, either through predicative primitives or a system of types.
- The selectional properties of words must be explained in terms of the lexical semantics of predication.
- An understanding of the semantics of nouns and the contribution that verbal arguments play in composition is crucial for an adequate model of how lexical semantics contributes to compositionality.

As we will see below, the first point makes an explicit connection between predicate decomposition theories (such as Lakoff, 1965/1970; Levin and Rappaport-Hovav, 1995) and type-theoretic approaches to lexical semantics (Dowty, 1979; Pustejovsky, 1995; Davis and Koenig, 2000; Asher and Pustejovsky, 2006; Asher, 2011). This in turn directly influences the manner in which selectional constraints are encoded. Finally, we will observe the central role of Aktionsarten and event typing in the determination of sentence meaning in composition.

There are essentially four strategies to lexical specification that have been adopted in the literature. These approaches can be defined in terms of how, or whether, they provide an intermediate level of interpretation from expressions in the object language to their denotations in the model. The interface defines the logical forms associated with lexical items in the language, and it is these expressions which are interpreted in the model.

The four approaches can be defined as follows:

Atomic predication: No interface language is provided, and there is direct interpretation of an object expression into the model.

Relational models: The interface is a relational structure, which is interpreted in the model.

Feature-based decomposition: Component-based features are used to classify an expression in the object language into distinct concepts in the model.

Structural decomposition: Component-based features are organized as a graph structure, with associated compositional interpretations in the model.

Atomic predication refers to non-decompositional approaches to lexical semantic interpretation, such as the early type-theoretic semantics of

Montague (1974), where the lexical items are primitives in the object language, or the “language of thought” proposed by Fodor (1975).

This view assumes a very tight coupling between the lexical semantics of a word and the syntactic projection associated with it. One consequence of this position is that there are as many lexical entries for a word as there are senses for it in the object language.

Relational models also start with the assumption that words are treated as primitives, but unlike atomic predication theory, they can have arbitrarily complex relational structures that facilitate semantic inferences. Such approaches are not as strict about introducing arguments that are visibly expressed in the syntax of the language, as is done with atomic predication models. Perhaps the best-known example of this strategy is Davidson’s addition of the event variable to action predicates (Davidson, 1967a), as well as most subsequent work assuming an event variable for eventuality predicates in language. Computational models for linguistic inference also often invoke this strategy in order to perform reasoning within established logical systems. Hobbs et al. (1993), for example, working within a framework of first-order abductive inference, adds any additional parameters to the argument structure of a predicate that are needed for an inference.

Feature-based decomposition has been used for discriminative analysis of natural language semantics since Nida (1949), and then, significantly in the 1960s, when the Katz and Fodor (1963) and Katz and Postal (1964) models within early generative grammar gained currency within the field. All expressions in the object language are decomposed into sets of binary-valued features, distinguishing concepts such as gender, number, age, marital status, and so on. Recently, some of these ideas have come back into vogue, where vector-based representations of word meaning have emerged as a way to handle some long-standing problems in the computational interpretation of language. These distributional semantic models utilize far more sophisticated techniques for identifying word distributions and computing similarity and relatedness with them (Schütze, 1993; Schütze, 1995; Landauer and Dumais, 1997; Padó and Lapata, 2007; Erk and Padó, 2008).

Finally, structural decomposition is the approach currently adopted by many lexical semanticists working in the interface between syntax and semantics. Words are defined as restricted algebraic structures, with primitive predicates as atomic elements. This approach has been adopted broadly, from Dowty (1979) in his model-theoretic interpretation of generative semantics (Lakoff, 1965/1970), to Van Valin Jr. (2005), Jackendoff’s Conceptual Structure (Jackendoff, 1983, 1990), and variants of this model (Levin and Rappaport-Hovav, 1995, 2005). Both Generative Lexicon Theory (Pustejovsky and Boguraev, 1993; Pustejovsky, 1995) and semantic interpretations for Head-Driven Phrase Structure Grammar (HPSG) (Davis and Koenig, 2000; Ginzburg and Sag, 2000) can also be seen as assuming a rich, structural decomposition as the foundation for their models.

In this entry, after a brief history of research in lexical semantics, we review the major issues confronting models of word meaning. Then we examine more closely the meaning components of individual words, beginning with word classes and types, and moving on to different strategies for encoding lexical meaning, focusing on polysemy and argument selection. Next we consider the role played by event semantics in the determination of meaning and survey approaches to decomposition and semantic roles. After a brief introduction to qualia structure, we conclude with an examination of the ways that word meanings contribute to the compositional mechanisms responsible for building sentence meanings.

2.2 The history of lexical semantics

Interest in the meaning of words goes back as far as philosophical thought and speculation about the natural world. Aristotle, for example, in his *De Interpretatione* (McKeon, 1968) outlines a broad theory of language meaning, framed within a context of determining truth in the service of reasoning. As part of his general characterization of nature, Aristotle proposes a classification of thoughts and the words that convey them using a systematic set of dimensions called *aitia*. Applied to verbs and their corresponding activities, this allows one to distinguish between *telic* (goal-completive) and *atelic* (non-goal completive) eventualities. This has both influenced work in lexical aspect and Aktionsarten within linguistics and provided much of the conceptual background for *qualia structure* (Pustejovsky, 1995), see Section 2.7.

With the exception of Locke (1690), Hume (1938), and Reid (1764), there was very little in the way of serious or systematic theorizing about word meaning until late in the nineteenth century. At this time, language researchers in Germany and France began focusing on lexical semantics from a psychological perspective, looking at the relation between words and concepts. Bréal (1897), for example, considered polysemy (a term he coined) to be a necessary creative component of language, claiming that this phenomenon better than most others in semantics illustrates the cognitive and conceptualizing force of the human species. Similarly, German semasiologists viewed the lexicon, word meaning, and mechanisms of polysemy as illustrative of the "life force" of human language (see Erdmann, 1900).

In the early twentieth century, European Structuralists introduced the distinction between syntagmatic and paradigmatic processes. Syntagmatic relations refer to the influence of "horizontal" elements on a word or phrase (co-occurrence), in contradistinction to paradigmatic processes, which refer to "vertical" or alternative substitutions in a phrase. The terms had significant currency in early and mid-twentieth-century linguistics from de Saussure on and helped to define the formal study of syntax as widely practiced today. The Structuralists saw paradigmatic relations encoded in the various lexical systems of a language, and this was elaborated into a framework of componential analysis for language meaning (Nida, 1949; Jakobson, 1973). Componential analysis, used by anthropologists to study kinship terminology, is a method for breaking down the meanings of lexical items into a set of features, thereby illustrating similarities and differences of meaning between the items. The goal of such analysis was simply to classify the lexical items in the language with some finite set of features, its ultimate contrastive elements. These contrastive elements are then structured in a matrix, allowing for dimensional analyses and generalizations to be made about lexical sets occupying the cells in the matrix. Hjelmslev (1943/1961), for example, decomposed lexical items into "paradigms" which pattern the same way distributionally in the language. The componential analysis of lexical items entails a decomposition into distinctive features: *man* as [+adult, +male], *woman* as [+adult, −male], *girl* as [−adult, −male], and so forth. Similar distinctions are made in Nida (1949) for distinguishing morphological patterns.

In the early days of generative linguistics, many of the ideas of the Structuralists found their way into the first formulations of lexical knowledge for transformational grammars. Of particular importance was Katz and Fodor's (1963) and Katz and Postal's (1964) theory of feature-based semantics, where the meanings of words were composed entirely of sets of features with Boolean values. In line with Chomsky (1957), the role of meaning was limited: to detect semantic anomalies and determine the number of readings associated with a sentence. The theory offered proposals for the decomposition of sentences and lexical items, with explicit rules for linking items to syntactic structure. While influential in the short term, this theory had no adequate theory of compositionality and was seen to be too weak as a model for natural language semantics (see Weinreich, 1972).

At the same time as features were being introduced into linguistic analyses, the role of a predicate's *valence* in relation to syntactic expressibility began to be studied. Valence, a term introduced by Tesnière (1959), is a characterization of the number and semantic nature of arguments carried by a verb or other predicative expressions. In the late 1960s and early 1970s, alternatives to the Katz and Fodor model began to emerge that incorporated many of the characteristics and principles of valence-based grammars. These theories attempt to respect the relational structure of sentence meaning while encoding the named "semantic functions" of the arguments in the lexical entries for predicates (Lakoff, 1965/1970; Fillmore, 1968a, 1969; Gruber, 1976; Jackendoff, 1972).

Fillmore (1968a, 1977a), for example, uses an enriched notion of valence to account for how arguments bind to syntactic positions in the sentence. From these early accounts of *case grammar*, Fillmore and colleagues developed a broader notion of *frame semantics* (cf. Fillmore, 1976, 1982), where human activities are conceptualized as lexically encoded frames. A semantic frame specifies the conditions under which a predicate combines with its possible arguments, seen as the participants in the event which that predicate denotes.

Some of these ideas were incorporated into lexically rich, feature-based semantics, in an attempt to explain how the semantic properties of predicates predict syntactic expressibility and behavior. One version of this grew into the framework known as *generative semantics* (Lakoff, 1965/1970; McCawley, 1968a), where the input to semantic interpretation was the deep structure of a sentence. While this started as an attempt to explain the selectional preferences imposed by verbs on their arguments, the scope of the theory expanded to account for all semantic interpretation from deep structure.

This view changed with Chomsky's and Jackendoff's *lexicalist* work in the 1970s, where the role of the lexicon became more central to grammatical processes, and generalizations could be made in terms of what properties were shared by lexical items (Chomsky, 1970; Jackendoff, 1972). While the *Aspects*-model of selectional features restricted the relation of selection to that between lexical items, work by McCawley (1968a) and Jackendoff (1972) showed that selectional restrictions must be available to computations at the level of derived semantic representation rather than at deep structure. Later work by Bresnan (1982), Gazdar et al. (1985), and Pollard and Sag (1994) extended the range of phenomena that can be handled by the projection and exploitation of lexically derived information in the grammar. In these frameworks, the lexicon plays a central part in the way compositional processes are carried out in the grammar.

Before the mid-twentieth century, there was little interest in word meaning within traditional philosophy. While linguistic semantics can trace its roots back to both Frege (1892b) and Russell (1905), these and other authors were less interested in word meaning and linguistic behavior than they were in how words were used as the medium through which judgments can be formed and inferences made. Frege's focus lay in formulating the rules which create meaningful expressions in a compositional manner, while also introducing an important distinction between an expression's sense and its reference. However, because of the descriptive bias in linguistics at the time, linguists largely failed to appreciate the role that systematic models of compositionality might play in language generally. Not until mid-century, with Montague's synthesis of grammatical description and intensional type theory, were these issues addressed in a comprehensive (and influential) manner.

Montague (1970a,b) introduces a bold new program for semantic interpretation in natural language, based on formal logic with a model-theoretic interpretation. Some of the most influential contemporary work on lexical semantics is based on this foundation, as we shall see in the sections below. In Dowty (1979), a model-theoretic interpretation of the decompositional techniques introduced by Lakoff, McCawley, and Ross was developed. Dowty's work, together with Partee (1975)'s seminal article on how

Montague semantics was compatible with syntactic models familiar to linguists at the time, helped introduce formally grounded approaches to semantics to the mainstream of the field.

Over the past two decades, there have been serious efforts to create a synthesis of lexical semantics and compositional mechanisms within linguistic theory. What has emerged is a new view of the role played by the lexicon regarding how composition is directed. Further, there is a new understanding of how contextualized aspects of interpretation impact the design decisions for what semantics is attributed to lexical items.

Examples of this approach can be seen in varieties of Generative Lexicon Theory (Pustejovsky, 1995; Bouillon and Busa, 2001), and in work by Jackendoff (1997, 2002). These developments have helped to characterize the approaches to lexical design in terms of a hierarchy of semantic expressiveness. There are at least three such classes of lexical description, defined as follows: *Sense Enumerative Lexicons*, where lexical items have a single type and meaning, and ambiguity is treated by multiple listings of words; *Polymorphic Lexicons*, where lexical items are active objects, contributing to the determination of meaning in context, under well-defined constraints; and *Unrestricted Sense Lexicons*, where the meanings of lexical items are determined mostly by context and conventional use. It seems clear that the most promising directions seem to be a careful and formal elucidation of polymorphic lexicons, and this will form the basis of our subsequent discussion.

2.3 Issues in lexical semantics

In this section, we look more closely at a set of phenomena in language that are of particular relevance to how lexical meaning is encoded and applied through compositional mechanisms in the grammar. We focus on four specific problems here: (a) the semantic distinctions between lexical ambiguity and systematic polysemy; (b) the nature of selectional preferences in verb-argument composition; (c) the polymorphic behavior of argument selection by predicates; and (d) the question of lexical representation, more generally construed.

2.3.1 Ambiguity versus polysemy

Given the compactness of a lexicon relative to the number of objects and relations referred to in the world, lexical ambiguity would seem inevitable for any language. Furthermore, the cultural, historical, and linguistic blending that contributes to the meanings of our words tends to make lexical ambiguity – when a word carries more than one meaning – appear arbitrary as well. Hence, *homonymy* – where one lexical form has many meanings – is to be expected in a language. Examples of homonyms are illustrated in the sentences below:

(1) a. Mary strolled along the *bank* of the river.
b. This company is the largest *bank* in the city.

(2) a. The doctor *treated* the woman with antibiotics.
b. He always *treats* his guests well.

(3) a. First we leave the gate, then we *taxi* down the runway.
b. John saw the *taxi* on the street.

(4) a. The judge asked the defendant to approach the *bar*.
b. The defendant was in the pub at the *bar*.

The senses of each noun in the examples above are arguably unrelated to each other; such lexical distinctions have also been called *contrastive ambiguities* (cf. Weinreich, 1972). For this reason, it is assumed that homonyms are represented as separate lexical entries within the organization of the lexicon. Words with multiple senses are simply listed separately in the lexicon, but this does not seem to compromise or complicate the compositional process of how words combine in the interpretation of a sentence.

We can compare this to the phenomenon known as *polysemy* (cf. Apresjan, 1973). Polysemy is the relationship that exists between related senses of a word, rather than arbitrary ones, as in the above examples. For example, the noun *book* is polysemous and can refer to either a physical object or the information contained in it, as illustrated in (5) below.

(5) a. Mary carried *the book* home.
b. Mary doesn't agree with *the book*.

Unlike the homonyms above, these two senses are logically related to each other by the very concept of *book*. Similarly, the noun *lunch* can refer to the food intended for consumption at a meal or the actual meal itself, as seen in (6).

(6) a. Mary has *her lunch* in her backpack.
b. The phone rang during *my lunch* yesterday.

In (7), a similar logical relation exists between the two senses of *flight*; in (7a), it refers to the event of flying, while in (7b) it refers to the plane engaged in flying.

(7) a. *The flight* lasted three hours.
b. *The flight* landed on time in Los Angeles.

While the two senses of the noun *bank* in (1), as well as of the other nouns in examples (2)–(4), are not related (except by phonological form), each of these examples indicates a formal relation between the lexical senses. It is the role of the lexicon to distinguish such ambiguities, and to establish what this logical relation is.

Polysemy is not an isolated phenomenon in language, but rather is associated with every major category in the lexicon, within all languages. For example, adjectives such as *good*, *dangerous*, and *fast* can be viewed as polysemous, where the sense is modulated slightly, depending on the noun being modified.

(8) a. John is a *a good teacher*.
    b. *A good meal* is what we need now.
    c. Mary took *a good umbrella* with her into the rain.

In each of these sentences, *good* is a manner modifier whose interpretation is dependent on the noun it modifies; in (8a) it means "to teach well", in (8b), it means a "tasty meal", and in (8c), it means "something keeping you dry". Similar remarks hold for the adjective *dangerous*.

(9) a. This is *a dangerous road* at night.
    b. She used *a dangerous knife* for the turkey.

That is, the road is dangerous in (9a) when "one drives on it", and the knife is dangerous in (9b) when "one cuts with it". Finally, the adjective *fast* in the sentences below acts as though it is an adverb, modifying an activity implicit in the noun; that is, *programming* in (10a) and *driving* in (10b).

(10) a. Mary is *the fastest programmer* we have on staff.
     b. The turnpike is *a faster road* than Main Street.

The exact nature of how adjectives are interpreted relative to the head is discussed in Pustejovsky (1995), and Busa (1996).

A somewhat related phenomenon involving adjectival scope instead of true polysemy can be seen in the sentences in (11), where the adjective can modify the agentive nominal activity or the nominal as an individual.

(11) a. Mary is *a beautiful dancer*.
     b. Jones is *a good judge*.

This problem has been studied by several researchers (Landman, 1989b; Pustejovsky, 1995; Busa, 1996; Larson, 1998), and while not involving polysemy, it does highlight the need for a potentially richer semantic content of agentive nominals than is typically assumed.

As mentioned above, polysemy is a phenomenon exhibited by all the major categories in language, including verbs. Because verbal polysemy is arguably the most common form of systematic sense relatedness, we cover it in more detail below.

What is important to realize from the data shown here, for both nouns and adjectives, is that simply listing the senses of these words will not always account for their creative use in novel contexts in language. Accounting for the behavior of polysemy in language is one of the most difficult problems facing a theory of the lexicon, and we take up the solutions to this problem in Section 2.8.

2.3.2 Selectional constraints

To demonstrate the interaction of the lexicon with the other components of grammar, let us examine how lexical information is accessed and exploited by syntactic and semantic operations, a process generally referred to as *selection*.

By far, the most widely studied type of selection process involves the constraints imposed on neighboring phrases by a word, by virtue of its lexical properties. One of the most important properties of a verb is an encoding of what phrases it can appear with in the language. In general, this is the problem of determining how many arguments a verb can appear with in the syntax and is referred to as *argument selection*. There is a general rule of thumb that the number of arguments that the predicate allows in the syntax corresponds to the number of participants an event has.[1] That is, the argument structure of a word loosely reflects the underlying relational meaning associated with that word.

For example, consider the behavior of one argument (intransitive), two argument (transitive), and three argument (ditransitive) verbs in English. The *arity* (number of arguments) information is encoded in the verb's *argument structure*. The verbs *laugh*, *see*, and *give* are simple examples in English.

(12) a. The man *laughed*.
b. The girl *saw* a bird.
c. The boy *gave* a treat to the dog.

The argument structure for the verbs in these sentences can be represented as follows:

(13) a. **laugh**(arg_1)
b. **see**(arg_1,arg_2)
c. **give**(arg_1,arg_2,arg_3)

The lexicon plays an important role in determining whether a linguistic expression in a language is well-formed, and selection is the mechanism through which this is accomplished. For example, it is because of the argument structures in (13) that these verbs do not appear in the wrong grammatical contexts, such as (14) below.

(14) a. *The man *laughed* the ball.
b. *A bird *saw*.
c. *The boy *gave* to the dog.

One of the most important themes when studying the lexicon is to establish just what role lexical information plays in determining whether an expression is well-formed or not.

[1] We will see below that this is not always the case, however, and there are more complex relations between semantic representations and the syntactic structures that may appear.

Although the argument structure indicates how many arguments to expect with a verb in the syntax, it says nothing about what *kind* of arguments these should be. This can be accomplished by adding *selectional constraints* to the arguments of the verb. A selectional constraint is a requirement on an argument that must be satisfied for the verb to be interpreted properly. The most important constraint imposed on an argument by a predicate is its *syntactic category*, that is, whether it is a Noun Phrase (NP), Verb Phrase (VP), Adjective Phrase (AP), or sentence (S). For example, both *write* and *think* are transitive verbs and have the argument structure in (15):

(15) V(arg_1,arg_2)

However, their arguments have different syntactic requirements, as we can easily see in (16).

(16) a. The girl *bought* [$_{NP}$ some books].
b. The girl *thinks* [$_{S}$ her brother is funny].

This can be encoded as part of the argument structure for that verb if we make a list of features for each argument. In our representation, this list will be shown directly following the argument it is associated with. So, for example, let `cat` be the feature for the category of the argument; then we have the following distinct argument structures for these two verbs.

(17) a. **buy**(arg_1[cat=NP], arg_2[cat=NP])
b. **think**(arg_1[cat=NP], arg_2[cat=S])

Lexical information impacting the grammar Thus far, the lexical information needed to interface to the syntax is of two sorts:

- *Argument structure* of a verb: this determines the number of phrases in the syntactic construction associated with the verb.
- *Syntactic category* of each argument: this identifies the actual syntactic phrase associated with the argument.

There are other selectional phenomena captured by constraints. For example, observe the differences in the sentences below.

(18) a. The man / *the rock laughed.
b. The man / the rock fell.

The constraints imposed on the subject by a verb such as *laugh* would include the feature of `animacy`, while a verb like *fall* does not. Since this is a lexical property of these verbs, their argument structures should reflect this distinction. Hence, the verb *laugh* has a selectional constraint indicating that the subject must be animate, as shown in (19).

(19) **laugh**(arg_1[cat=NP,animacy=+])

Similarly, the feature of `number` is important for selecting the appropriate subject for verbs such as *assemble*, *gather*, and *disperse*; namely, the subject must be plural.

(20) a. The students *gathered* in the courtyard.
b. The children *assembled* in the gym.
c. The fans *dispersed* from the stadium.

Notice that it can be either *grammatically* plural, as in (20) above, or *semantically* plural, as with the subjects in (21).

(21) a. The class *gathered* in the courtyard.
b. The team *assembled* in the gym.
c. The crowd *dispersed* from the stadium.

That is, the nouns *class*, *team*, and *crowd* can all be considered as non-atomic, from the point of view of the verb and the grammar. Hence, the argument structure for these verbs must reflect this phenomenon by adding a selectional constraint referring to `number`, ranging over the values “singular” and “plural”.

(22) **assemble**(arg_1[cat=NP,num=plural])

It is interesting to ask whether these selectional constraints are independent or whether they interact with one another. Let us look a little more carefully at the verbs in (20) and (21) above. We saw that the subject for each of these verbs must be plural. How does the feature `number` interact with `animacy`? Notice that both *disperse* and *gather* allow non-human subjects, while *assemble* does not.

(23) a. The clouds *gathered* in the sky.
b. *The clouds *assembled*.
c. The clouds *dispersed* quickly.

To ensure that a structure such as (23b) is not licensed by the grammar, there must be an encoding of how the structural conditions are being violated. Notice that this can already be done using the constraints we’ve discussed, namely the combination of `number` and `animacy`, as illustrated in (24) below.

(24) **assemble**(arg_1[cat=NP,num=plural,animacy=+])

It would appear, then, that the verb *assemble* requires the plural subject to be animate, while the very similar verbs *gather* and *disperse* do not.

Selectional constraints, in fact, can determine the acceptability of arguments in any positions in the grammar; consider the distinction between the verbs *force* and *convince*. Although they are synonyms in many contexts, *force* and *convince* have different selectional properties, the latter requiring that the “convinced” be a cognitive agent.

(25) a. John *forced* the guest to leave.
b. John *convinced* the guest to leave.

(26) a. John *forced* the door to open.
b. *John *convinced* the door to open.

Verbs select for different spatial prepositions as well; verbs involving spatial manner descriptions, for example, are quite specific for the kinds of prepositions they allow. Consider, for example, *on* and *in* versus *over* in the sentences below, when selected by the verb *lie*.

(27) The cat is lying on the floor / in the box / *over the floor / . . .

Selectional constraints can also be enlisted to determine the appropriate selection of *manner adverbials* for verbs, as well. Notice that in (28), without some sort of adverbial phrase associated with the verbs *behave* and *perform*, the sentences below are ungrammatical.[2]

(28) a. Mary *behaved* *(badly).
b. John *performed* *(admirably).

We see from this brief discussion that selection is an important part of how lexical information is conveyed to the syntactic operations in the grammar.

Lexical information impacting the grammar From what we discussed, there are at least three types of lexical information which interface to the syntax:

- *Argument structure* of a verb: this determines the number of phrases in the syntactic construction associated with the verb.
- *Syntactic category* of each argument: this identifies the actual syntactic phrase associated with the argument.
- *Selectional constraints* of each argument: these identify specific grammatical and semantic features of the argument being selected by the verb. Selectional constraints have been used to encode all manner of lexical and syntactic restrictions in grammar. As such, they are part of the meaning of words and have direct effects on the syntax, morphology, and semantics of the language.

2.3.3 Verb meaning and mapping to syntax

We saw in the previous section the role that the lexicon plays in ensuring that a linguistic expression is well-formed. Recall that it is the argument structure that prevents verbs from ending up in the wrong syntactic contexts in sentences (Levin, 1993). That is, if the lexicon says it is so, then the grammar follows.

[2] It should be noted that (28a) does have a default reading of “behave well”, when no adverb is present.

(29) a. **laugh**(arg_1) $\Longrightarrow$ The man *laughed*.
b. **see**(arg_1,arg_2) $\Longrightarrow$ The girl *saw* a bird.
c. **give**(arg_1,arg_2,arg_3) $\Longrightarrow$ The boy *gave* a treat to the dog.

Or does it? Notice that there are contexts for each of the verbs in (29), which exploit properties of the verb, giving rise to constructions not allowed by the original argument structure.

(30) a. The man *laughed* himself sick.
b. The girl *saw* a bird fly into the room.
c. The man *gave* at the office.

In each of these sentences, the argument structure for the verb has been violated in some way. In (30a), there is an additional NP object to the verb *laugh*, and a predicate modifying the object NP *himself*. In (30b), what the girl saw was not just a bird, but what the bird did; namely, “flying into the room”, described as an additional VP. Finally, in (30c), both of the expected arguments are missing, and presumably inferred from the context of the utterance.

These illustrate one aspect of the phenomenon of *verbal polysemy*. Polysemy, as defined in the previous section, is the term given to an ambiguity where the different meanings of the word are logically related to each other. Many verbs can appear in multiple contexts taking a different number of arguments in each, a phenomenon known as an *alternation*. For example, the verbs *break*, *roll*, and *sink* all have intransitive and transitive forms, as shown in the sentences below.

(31) a. The glass *broke*.
b. Mary *broke* the glass.

(32) a. The ball *rolled* down the slide.
b. The boy *rolled* the ball down the slide.

(33) a. The ship *sank*.
b. The torpedo *sank* the ship.

How does the lexicon represent such ambiguities? The simplest way would be to list the different argument structures for each verb, as shown in (34).

(34) a. $\textbf{break}_1$(arg_1); $\textbf{break}_2$(arg_1,arg_2)
b. $\textbf{roll}_1$(arg_1); $\textbf{roll}_2$(arg_1,arg_2)
c. $\textbf{sink}_1$(arg_1); $\textbf{sink}_2$(arg_1,arg_2)

Note that the semantic role of the intransitive subject is the same as the transitive object NP in each one of these verbs. That is, it is the *undergoer* or *patient*.

This kind of alternation does not apply to all intransitive verbs, of course, and the lexicon must somehow prevent verbs like *arrive* and *die* from becoming transitive.

(35) a. Your package *arrived*.
b. *The mailman *arrived* your package.

(36) a. My computer *died* yesterday.
b. *The storm *died* my computer yesterday.

The question arising from such cases is this: what allows for an alternation to occur for some verbs while not for others? Is this part of the lexicon or some other rule or strategy in the grammar?

Verb alternations are quite prevalent in language and pose a difficult problem to lexicon designers. For example, how many different argument structures do we have for a verb like *sweep*?

(37) a. John *swept*.
b. John *swept* the floor.
c. John *swept* the dirt into the corner.
d. John *swept* the dirt off the sidewalk.
e. John *swept* the floor clean.
f. John *swept* the dirt into a pile.

Such alternating patterns typically apply to more than just one verb; hence, it would be more efficient to have a general strategy for how these word senses are related rather than simply listing the different senses for each verb.

Lexical information is sensitive to both the syntax and the discourse context. Although a verb may be lexically specified to have a certain number of arguments, there are many situations in which this can be violated. We saw this above in (30c) with the verb *give*. In fact, this is quite common but appears to be governed by systematic rules, and not just pragmatic information. For example, while (30c) is grammatical, for most speakers (38b) below is not.

(38) a. The woman *donated* her car to the foundation.
b. *The woman *donated*.

Similarly, *eat* and *drink* can appear without their direct objects, but *devour* and *gulp* cannot (at least for most speakers).

(39) a. The girl *ate* her lunch quickly.
b. The girl *ate* quickly.

(40) a. The dog *devoured* the bone.
b. *The dog *devoured*.

(41) a. The boy *drank* his milk quietly.
b. The boy *drank* quietly.

(42) a. The boy *gulped* his milk quietly.
b. *The boy *gulped* quietly.

This type of alternation is typically called *indefinite NP deletion* and is related to other "pragmatic deletions", such as those shown below.

(43) a. John *tried* to call his mother yesterday.
b. John *tried* yesterday.

(44) a. John *attempted* to call his mother yesterday.
b. *John *attempted* yesterday.

The ability to ignore part of the argument structure in the syntax seems to be a lexical property, one which is idiosyncratic to each verb.

Finally, we consider a kind of verbal polysemy not involving argument alternation, but a *syntactic category alternation*. Recall from our discussion above (see (17)) that we motivated a way to distinguish between different syntactic types for an argument; that is, *buy* and *think* have different category values associated with their second arguments:

(45) a. **buy**(arg_1[cat=NP],arg_2[cat=NP])
b. **think**(arg_1[cat=NP],arg_2[cat=S])

What happens, however, when the syntactic distinction involves the same verb? Consider the sentences below, where the verb *begin* appears in three distinct syntactic contexts.

(46) a. Mary *began* to read the novel.
b. Mary *began* reading the novel.
c. Mary *began* the novel.

Verbs like *hate* and *love* in English can appear in even more contexts, as seen in (47) below.

(47) a. John would *hate* Bill to leave.
b. John *hates* (it) that Bill left.
c. John *hated* to lose the game.
d. John *hated* losing the game.
e. John *hated* that he lost the game.

The examples above in (46) and (47) bring up the issue of how to make generalizations in the lexicon; that is, is there a "compact" manner in which to express that the verb *begin* means the same thing in each sentence in (46), and likewise for the verb *hate* in (47).

2.4 Event semantics

In this section, we look at the notion of event in lexical semantics. There are two traditions to examine when studying the role of events in the semantics of language:

Event typology: Predicates in language have an *event variable* that can be treated as a first-order individual in the semantics, to enable logical inference (Davidson, 1967a).

Aktionsarten: Predicates in language can be classified according to their event type or aspectual class, in order to specifically capture grammatical and semantic behaviors (Vendler, 1957).

The move by Davidson to introduce a first-order event variable in the representation was mainly motivated by the need to provide a coherent analysis of adverbial modification in the interpretation of sentences. Under this proposal, two-place predicates such as *eat* and three-place predicates such as *give* contain an additional argument, the event variable, e, as depicted below.

(48) a. $\lambda y \lambda x \lambda e[\text{eat}(e, x, y)]$
b. $\lambda z \lambda y \lambda x \lambda e[\text{give}(e, x, y, z)]$

In this manner, Davidson is able to capture the appropriate entailments between propositions involving action and event expressions through the conventional mechanisms of logical entailment. For example, to capture the entailments between (49b–d) and (49a) below,

(49) a. Mary ate the soup.
b. Mary ate the soup with a spoon.
c. Mary ate the soup with a spoon in the kitchen.
d. Mary ate the soup with a spoon in the kitchen at 3:00 pm.

In this example, each more specifically described event entails the one above it by virtue of *and*-elimination (conjunctive generalization) on the expression.

(50) a. $\exists e[\text{eat}(e, m, \text{the\_soup})]$
b. $\exists e[\text{eat}(e, m, \text{the\_soup}) \wedge \text{with}(e, \text{a\_spoon})]$
c. $\exists e[\text{eat}(e, m, \text{the\_soup}) \wedge \text{with}(e, \text{a\_spoon}) \wedge \text{in}(e, \text{the\_kitchen})]$
d. $\exists e[\text{eat}(e, m, \text{the\_soup}) \wedge \text{with}(e, \text{a\_spoon}) \wedge \text{in}(e, \text{the\_kitchen}) \wedge \text{at}(e, \text{3:00 pm})]$

There are of course many variants of the introduction of events into predicative forms, including the identification of arguments with specific named roles (or partial functions, see Chierchia, 1989; Dowty, 1989).

In lexical semantic analysis, it is standard practice to create component-based classifications using linguistic data that demonstrate pairwise distinctions for grammatical or semantic well-formedness judgments. One such approach is the determination of aspectual class or *Aktionsart* (see also Rothstein, Chapter 12). This is essentially a characterization of the different kinds of eventualities that predicative expressions denote. There have been several influential distinctions proposed in the literature, but the best known are those introduced by Kenny (1963) and Vendler (1967). Kenny assumed that there are three basic aspectual types: *states*, *activities*, and *performances*. Vendler proposes a similar distinction for states and processes, but splits the last class (his *events*) into two categories, *accomplishments* and *achievements*. His classification as well as his terminology have been the starting point for much of the work in aspect and event semantics in the field. These event classes are summarized briefly below.

State: An unchanging situation that holds at a time or over an interval, with no endpoints.
Activity: A dynamic event with no endpoints (an *atelic* event).
Accomplishment: An incrementally changing event with a culminating endpoint (a *telic* event).
Achievement: An instantaneous change of state.

Examples of states are seen in simple attributive predications, such as (51a–b), as well as with non-dynamic relations, such as (51c–d).

(51) a. Mary is *happy*.
b. Marc is *Dutch*.
c. Jan *loves* Bill.
d. Mary *believes* Jan is happy.

States can be distinguished from activities by virtue of certain grammatical diagnostics, many of which Vendler introduced into the literature. Consider a verb such as *swim* in sentence (52a). This denotes an activity of unspecified duration, and the sentence does not convey information regarding a culmination of this activity. With the addition of a goal *to*-PP, however, as in sentence (52b), the swimming activity is bounded, and the resulting event denotes an accomplishment.

(52) a. Jan *swam* yesterday.
b. Jan *swam* to the dock yesterday.
c. Jan *is swimming*.

Hence, we can analyze the verb *swim* as lexically denoting an activity, which can be interpreted contextually as an accomplishment, through the appropriate compositional modification. Finally, notice that sentence (52c) denotes a "snapshot" of the swimming process. Both of these aspectual *coercions* are known to distinguish processes from states.

In a similar fashion, there are verbs which lexically denote accomplishments, such as *build* and *destroy*. These verbs encode the logical culmination to the activity performed during the event, as illustrated in (53) below.

(53) a. The children *built* a fort.
b. John *destroyed* the anthill.

Finally, an achievement is an event that results in a change of state, just as an accomplishment does, but where the change is thought of as occurring instantaneously. For example, in sentences (54a–c) the change is not a gradual one, but something that has a point-like quality to it. Hence, modification by *point adverbials* such as *at 3 pm* is suggestive that a sentence denotes an achievement (see Dowty, 1979).

(54) a. The plane *crashed* at noon.
b. John *found* his wallet at 3 pm.
c. The train *arrived* at midnight.

Pustejovsky (1988, 1991b) extends the decompositional approach presented in Dowty (1979) by explicitly reifying the events and subevents in the predicative expressions. Unlike Dowty's treatment of lexical semantics, where the decompositional calculus builds on propositional or predicative units (as discussed above), a "syntax of event structure" makes explicit reference to quantified events as part of the word meaning. Pustejovsky further introduces a tree structure to represent the temporal ordering and dominance constraints on an event and its subevents (see also Moens and Steedman, 1988).

(55) a. EVENT → STATE | PROCESS | TRANSITION
b. STATE: → e
c. PROCESS: → $e_1 \ldots e_n$
d. TRANSITION$_{ach}$: → STATE STATE
e. TRANSITION$_{acc}$: → PROCESS STATE

For example, the accomplishment denoted by "building a house" consists of the building process, followed by the state representing the result of the object being built. Grimshaw (1990) adopts this theory in her work on argument structure, where complex events such as *break* are given a similar representation. In such structures, the process consists of what an agent, x, does to cause the breaking of y, and the following state represents the resulting broken item. The process corresponds to the outer causing event as discussed above, and the state corresponds in part to the inner change-of-state event. Both Pustejovsky and Grimshaw differ from the authors above in assuming a specific level of representation for event structure, distinct from the representation of other lexical properties. Furthermore, they follow Davidson (1967a), Higginbotham (1985), and particularly Parsons (1990), in adopting an explicit reference to (at least) one event variable in the parameter structure of the verbal semantics.

Recently, Rappaport-Hovav and Levin (2001) and Levin and Rappaport-Hovav (2005) have adopted a large component of the event structure model for their analysis of the resultative construction in English; event decomposition has also been employed for properties of adjectival selection, the interpretation of compounds, and stage and individual-level predication (Carlson, 1977a; Diesing, 1992a; Jäger, 2001).

Research done by Tenny (1989), Dowty (1991), Krifka (1992, 1998), and others enriches this typology by developing a theory of how an event is shaped by the incremental participation of the theme in that event. The central aspect of this view is that an accomplishment is defined as involving a homomorphism from parts of the event to parts of the incremental theme. The problem of incrementality can be illustrated with the following examples.

(56) a. John ate a hamburger.
b. Mary wrote a novel.

The process of eating something is an incremental activity and results in an accomplishment, described by reference to the quantized unit appearing in the direct object (theme) position, only when the entire hamburger has been eaten or the entire novel has been written.

Recent work on scalar change (see Hay et al., 1999; Beavers, 2008; Kennedy and Levin, 2008) and dynamic event semantics (Naumann, 2001) suggests a new understanding of the interplay between verb meaning, event semantics, and argument structure with these predicates, by focusing on the measurement of the change in value over the properties of the participants in each intermediate state during the event.

Starting from a quite different perspective, Krifka (1989a, 1992) presents an interpretation of Aktionsart using a lattice-theoretic interpretation of event structure. Using the sum operation from lattice theory, where a complex event $e_1 \sqcup e_2$ can be formed from any two events, e_1 and e_2, Krifka introduces a part/whole relation, $e_1 \sqsubset e_2$, iff $e_2 = e_1 \sqcup e_2$. This allows for a distinction to be made between *quantized* and *cumulative* predicates. Processes of undifferentiated activity, such as *walking* or *singing*, are cumulative and are closed under the sum operator, while an accomplishment, such as *build a house*, is not.

(57) a. $\forall P[\text{CUM}(P) \leftrightarrow \forall x \forall y[[P(x) \wedge P(y)] \rightarrow P(x \sqcup y)]]$
b. $\forall P[\text{QUA}(P) \leftrightarrow \forall x \forall y[[P(x) \wedge P(y)] \rightarrow x \not\sqsubset y]]$

These two classes map fairly closely to the well-known categories of atelic and telic predicates, mentioned above. For example, no part of building a house is a complete house-building; the event is quantized by the relationship to the result, i.e., the house (see Krifka, 1992). The activity of walking, on the other hand, is cumulative, and any event of walking is composed of subevents of walking. See Champollion and Krifka, Chapter 13 for further discussion.

2.5 Lexical decomposition

Lexical decomposition is concerned with the internal semantic structure of lexical items within a lexicon. The focus of lexical decomposition is on how the lexical items are semantically similar and distinct by virtue of shared knowledge structures or semantic primitives. While numerous structures and primitives (e.g., semantic features and semantic markers) have been proposed, we limit our discussion here to the most currently relevant theories, all of them *structural decomposition* models, as introduced above.

The goal of lexical decomposition has traditionally been to provide the necessary and sufficient conditions for the meaning of every lexical item in a subject domain or language. In many ways, this goal is similar to that of the syntactic analysis of sentences in a language. If primitives and structures are taken as an exhaustive set on top of which all expressions in the language are expressed, then the meaning of any lexical item in the language will have to be derived from these terms.

Perhaps the first significant framework of structural decomposition was generative semantics (see Lakoff, 1965/1970). It emerged just as Katz and Fodor's feature-based decomposition model was shown to be both incomplete and inadequate (see Weinreich, 1972) as a model for language semantics. Generative semanticists argued that deep structure was the appropriate level for semantic structure and that transformational rules operate only on that structure, without subsequent semantic consequences. Lexical items are decomposed into a set of abstract components and transformations are done on those components (see Lakoff, 1965/1970; Ross, 1970). For example, the lexical item *kill* is decomposed into the predicate DEAD and two higher-level predicates CAUSE and BECOME. The deep semantic structure associated with this verb consists of these semantic components organized into a labeled bracketing.

While the framework of generative semantics is no longer generally adopted, many aspects of the theory can be seen in contemporary lexical theories, owing to some significant generalizations the theory made about semantic structure in language. For example, the sentences in (58) form a sort of paradigm for the concept *dead*, related by the application of abstract predicates, which systematically change the meaning of the words associated with it (see Lakoff, 1965/1970; McCawley, 1968b; Carter et al., 1988).

(58) a. John killed Bill.
b. Bill died.
c. Bill is dead.

Assuming that the underlying form for a verb like *kill* encodes the stative predicate in (58c) and the relation of causation, generative semanticists posited representations such as (59) below.

(59) (CAUSE(x, (BECOME(NOT(ALIVE y)))))

Here the predicate CAUSE is represented as a relation between an individual causer x and an expression involving a change of state in the argument y. Although there is an intuition that the cause relation involves a causer and an event, neither Lakoff nor Carter makes this commitment explicitly (a point we return to below).

In a very influential work, Dowty (1979) presents a model-theoretic interpretation of the ideas from generative semantics, but with one important difference: Dowty associates complex decompositional expressions such as the one in (59) with the lexical items directly, without the syntactic machinery of predicate reduction that was required in a generative semantics derivation from deep to surface structure. Hence, the lexicon in Dowty's theory is far richer than normally envisioned.

Another significant contribution made by Dowty's study of word meaning is his reinterpretation of Vendler's classification of the Aktionsarten, mentioned in Section 2.4, in model-theoretic terms. Thus, states, activities, accomplishments, and achievements are given a formal treatment within an intensional logic.

Pursuing a similar line of reasoning to the generative semanticists concerning lexical decomposition, Jackendoff (1972) built on the Gruber (1976) thesis to argue that predicates such as [+cause] and [+inchoative] are encoded in the meaning of the word itself.[3] In Jackendoff's later work (1983, 1990, 2002), the approach to lexical decomposition makes claims for cognitive relevance that are not important motivations for many researchers. However, Jackendoff believes in the cognitive primacy of the primitives used within his system, and the role of these primitives in performing inferences. We examine Jackendoff's semantic representation briefly here, as it pertains to lexical decomposition.

At the core of Jackendoff's conceptual semantics is an inventory of primitive categories from which lexical meanings are constructed. These include concepts such as event, state, action, place, path, property, and amount. Along with these are the formation rules used to combine these categories into concepts associated with lexemes. To illustrate, consider sentence (60) below, along with its lexical representation in Jackendoff's system.

(60) a. John flew the plane from New York to Boston.
b. CAUSE(John, GO(plane, $[_{[\text{path}]}$ FROM (IN [Boston]), TO(IN [New York])]))

Interestingly, many of these predicates are spatial in origin. Jackendoff, like many others interested in the nature of semantic extension, explores the idea that the predicates within the spatial domain may be used to analyze concepts within other semantic domains. The basic idea here is that for different semantic fields such as possession, identification, circumstantial, and temporal, a verb from the spatial field can acquire a new meaning using the same primitives because it is being evaluated relative to a new field. Examples of this include:

(61) a. *Possessional*: John will take the house and Mary the kids.
b. *Identificational*: John turned into a great father.
c. *Circumstantial*: Mary led me to believe she was younger than I.
d. *Temporal*: The seminar has been moved to Wednesdays.

[3] Jackendoff (1972), following Chomsky (1970), was philosophically at odds with the claims made by generative semantics. The interpretive semantics they defended took semantics to be a process of interpretation over the derived surface form of the sentence, rather than the deep structure.

Regardless of the decompositional approach adopted, there are many questions that remain: do its primitives provide a complete and exhaustive description of the concepts of natural language? How open-ended is the system? What is the additional information which distinguishes the meaning of one lexical item from another? That is, the primitives do allow for generalizations, but what allows for the discrimination of concepts (the function filled by distinguishers in Katz and Fodor's theory), and how many features are necessary to that end?

In the approaches described here, the general approach is to use decomposition as "annotations" that reflect components of the meanings of lexical items, rather than exhaustively decompose the meanings of words, as was the goal in the earlier days of generative grammar.

More recent versions of lexical representations inspired by generative semantics and Jackendoff's decomposition model can be seen in the Lexical Relational Structures of Hale and Keyser (1993), where syntactic tree structures are employed to capture the same elements of causation and change of state as in the representations of Carter, Jackendoff, and Dowty. The work of Levin and Rappaport, building on Jackendoff's Lexical Conceptual Structures, has likewise been influential in further articulating the internal structure of verb meanings (see Levin and Rappaport-Hovav, 1995).

Generative Lexicon can be seen as a hybrid theory, incorporating aspects of Jackendoff's and Dowty's work with the event semantics of Davidson and Parsons. It has taken several decades for Davidson's observations regarding the role of events in the determination of verb meaning to find their way convincingly into the major linguistic frameworks. Over the past two decades, a synthesis has emerged which attempts to model verb meanings as complex predicative structures with richer event structures (see Parsons, 1990; Pustejovsky, 1991b, 1995; Levin and Rappaport-Hovav, 1995). This research has developed the idea that the meaning of a verb can be analyzed into a structured representation of the event that the verb designates, and has furthermore contributed to the realization that verbs may have complex, internal event structures.

2.6 Semantic roles

While decomposition aims to define lexical meaning in terms of a word's internal features, theories of semantic roles can be seen as *partial decomposition* models, focusing on articulating the function or role that the arguments of a predicate play in the determination of sentence meaning. Hence, as mentioned above in Section 2.2, we can view semantic roles as enhanced argument structure specifications.

Argument structure is the syntactic encoding of the functional behavior of an expression in the object language. In linguistically motivated models of lexical semantics, one of the basic semantic distinctions between nouns and verbs is stated in terms of their selectional behavior in syntax (the valence properties discussed in Section 2.2). The argument structure for a word can be seen as the simplest specification of its semantics, indicating the number and type of parameters associated with the lexical item as a predicate. For example, the verb *sleep* can be represented as a predicate taking one argument, *love* as taking two arguments, while the verb *buy* takes two, three, or four arguments, depending on the context.

(62) a. **sleep**(x)
b. **love**(x,y)
c. **build**(x,y,z)

The first theory of semantic roles within a generative model was *Case Grammar*, a semantically oriented grammar developed by Fillmore (1968b) and others, e.g., Anderson (1977) and Starosta (1988). Case was first used for the morphological analysis of noun endings in, e.g., German and Russian. Fillmore showed that these noun endings serve the same purpose as the positioning of nouns and prepositions in lexical surface structures. Fillmore (1968b) introduced the notion of case grammar and a case frame, a predicate containing arguments that are a set of obligatory and optional cases. Implicit in the theory is that each NP in a sentence can be assigned only one case and that the cases assigned by a verb can be realized only once in a sentence. He defined a number of cases including:

1. Agentive (A): an animate perceived instigator of the action identified by the verb.
2. Instrumental (I): an inanimate force or object causally involved in the action or state identified by the verb.
3. Dative (D): an animate being affected by the action identified by the verb.
4. Objective (O): the semantically most neutral case.

For example, the verb *open* requires that its objective role be filled: something must open or be opened. In (63a), only the objective role is filled. In (63b) and (63c), both agentive and objective roles are filled. In (63d), the instrumental and objective roles are filled. In (63e), the agentive, objective, and instrumental roles are filled.

(63) a. The door opened.
b. Mary opened the door.
c. The door was opened by Mary.
d. The key opened the door.
e. Mary opened the door with the key.

Fillmore noted that different case roles can occupy the same grammatical function, e.g., the grammatical subject is *door* in (63a) and (63c) which occupies the objective role, *Mary* in (63b) and (63c) which has the agentive role, and *key* in (63d) and (63e) which has an instrumental role.

Fillmore's theory attempts to explain how the arguments of a predicate are assigned to particular syntactic structures and is not concerned with establishing an independent level of semantic representation. To handle assignment, Fillmore assumed a principle of a case hierarchy, which allowed the selection of grammatical subject and object in default configurations. As an example of case assignment, consider the verb *break* which assigns both obligatory and optional cases. Obligatory cases include the objective, while optional cases include both the agentive and the instrumental cases.

(64) John broke the window with a rock.

In sentence (64), by default assignment, agentive is assigned to the subject, objective to the object, and instrumental to the phrase within the prepositional phrase. This is a type of default selection which can be violated if the verbal morphology indicates a structure such as a passive as in (63c).

Fillmore (1977a) notes some deficiencies and limitations of case grammar. Although case grammar is useful as a guide to decomposition of verb meaning, the theory does nothing to clarify the nature of sense relations like antonymy and hyponymy. Case grammar allows the recognition of sentences as partially synonymous by matching case structures. Because it does not claim to be a complete model for meaning, case grammar avoids the pitfalls of the decomposition model of Katz–Fodor; on the other hand, because of its incompleteness, it is unclear what role such a system of cases might play in a comprehensive semantic theory.

Part of Jackendoff's early work on the theory of lexical decomposition (as discussed above) included a formalization of Gruber (1976)'s more limited notion of thematic role. In this view, which is quite similar to Fillmore's conception of cases, Jackendoff (1972) classifies verbs according to which thematic (*theta*) roles they assign to their arguments. Thus, *theta theory*, as this view has come to be known, is a minimal decomposition of verb meaning, where features and abstract predicates have been replaced by named functional roles, as in case grammar. The success of theta theory has come not from how well it characterizes verb denotations, but in how "functional" roles interact with other principles of syntax to determine the well-formedness of a sentence.

Thematic relations are now generally defined as partial semantic functions of the event being denoted by the verb or noun, and behave according to a predefined calculus of roles relations (e.g., Carlson, 1984; Chierchia, 1989; Dowty, 1989). For example, semantic roles such as agent, theme, and goal can be used to partially determine the meaning of a predicate when they are associated with the grammatical arguments to a verb.

(65) a. **put**⟨AGENT,THEME,LOCATION⟩
b. **borrow**⟨RECIPIENT,THEME,SOURCE⟩

Thematic roles can be ordered relative to each other in terms of an implicational hierarchy. For example, there is considerable use of a universal subject hierarchy such as shown below (see Fillmore, 1968b; Comrie, 1981).

(66) AGENT > RECIPIENT/BENEFACTIVE> THEME/PATIENT > INSTRUMENT > LOCATION

Many linguists have questioned the general explanatory coverage of thematic roles, however, and have chosen alternative methods for capturing the generalizations they promised. Dowty (1991) suggests that theta-role generalizations are best captured by entailments associated with the predicate itself. A theta-role can then be seen as the set of predicate entailments that are properties of a particular argument to the verb. Characteristic entailments might be thought of as prototype roles, or *proto-roles*; this allows for degrees or shades of meaning associated with the arguments to a predicate. Others have opted for a more semantically neutral set of labels to assign to the parameters of a relation, whether it is realized as a verb, noun, or adjective (see Hobbs et al., 1988).

In the late 1970s, Fillmore began to rethink his views of semantics and developed a "scenes-and-frames" semantics to overcome the deficiencies and limitations he had noted with case grammar. The basic ideas of the new semantics, influenced in part by Minsky (1975), were that "people associate certain scenes with certain linguistic frames" and that "meanings are relativized to scenes," i.e., lexical items or expressions are understood by placing them within scenes or images in which they have some linguistic functions, e.g., naming.

In the theory, frames are linguistic entities that represent the meanings of lexical items. Scenes can not only be visual but also refer to "interpersonal transactions, standard scenarios, familiar layouts" (Fillmore, 1977b) and "body image" (the way we orient and classify linguistic frames such as UP-DOWN and LEFT-RIGHT).

Fillmore's account views scenes-and-frames as offering an alternative to traditional accounts of lexical decomposition issues like lexical ambiguity, synonymy, semantic fields and selection restrictions (Fillmore, 1975). Synonymy can be understood as "indistinguishable scenes for which the frame offers lexical options". A selection restriction is viewed as the relation between a frame and a scene. "The selection restriction information about the use of a word can be stated as a specification of the nature of the appropriate scene" (Fillmore, 1975, p. 71).

Fillmore furthermore argues that the denotations of lexical items can better be characterized by a sort of prototype theory; that is, "instead of the meaning of a linguistic form being represented in terms of a checklist of conditions that have to be satisfied in order for the form to be appropriately or truthfully used, it is held that the understanding of meaning requires, at least for a great many cases, an appeal to an exemplar or prototype" (Fillmore, 1975). A prototype is an ideal example of some concept to which other instances of the same concept bear a strong (but not necessarily total) resemblance. For example, a prototypical bird might have a beak, wings, feathers, and be able to fly, but an ostrich would still be recognized as a bird. The prototype idea can be seen in Berlin and Kay (1969)'s study of color terms, in Labov's work on the boundary criteria between cups and bowls (Labov, 1972b), and in the work of the psychologist Rosch and Mervis (1975).

2.7 Qualia structure

Thus far, we have focused on the lexical information associated with verb entries. All of the major categories, however, are encoded with syntactic and semantic feature structures that determine their constructional behavior and subsequent meaning at logical form. In Generative Lexicon Theory (Pustejovsky, 1991a, 1995), it is assumed that word meaning is structured on the basis of four generative factors, or *qualia roles*, that capture how humans understand objects and relations in the world and provide the minimal explanation for the linguistic behavior of lexical items (these are inspired in large part by the Moravcsik (1975, 1981) interpretation of Aristotelian *aitia*). These are: the FORMAL role: the basic category that distinguishes the object within a larger domain; the CONSTITUTIVE role: the relation between an object and its constituent parts; the TELIC role: its purpose and function; and the AGENTIVE role: factors involved in the object's origin or "coming into being". Qualia structure is at the core of the generative properties of the lexicon, since it provides a general strategy for creating new types. For example, consider the properties of nouns such as *rock* and *chair*. These nouns can be distinguished on the basis of semantic criteria which classify them in terms of general categories such as `natural_kind`, `artifact_object`. Although very useful, this is not sufficient to discriminate semantic types in a way that also accounts for their grammatical behavior. A crucial distinction between *rock* and *chair* concerns the properties which differentiate *natural kinds* from *artifacts*: functionality plays a crucial role in the process of individuation of artifacts, but not of natural kinds. This is reflected in grammatical behavior, whereby *a good chair*, or *enjoy the chair* are well-formed expressions reflecting the specific purpose for which an artifact is designed, but *good rock* or *enjoy a rock* are odd out of context, since for *rock* the functionality (i.e., TELIC) is undefined. Exceptions exist when new concepts are referred to, such as when the object is construed relative to a specific activity, such as in *The climber enjoyed that rock*; *rock* itself takes on a new meaning, by virtue of having telicity associated with it, and this is accomplished by integration with the semantics of the subject NP. Although *chair* and *rock* are both `physical_object`, they differ in their mode of coming into being (i.e., AGENTIVE): artifacts are manufactured, *rocks* develop in nature. Similarly, a concept such as *food* or *cookie* has a physical manifestation or denotation, but also a functional grounding, pertaining to the activity of "eating". These distinct aspects of a category are represented by the qualia structure for that concept, which provides a coherent structuring for different dimensions of meaning.

By analyzing the semantics of objects in terms of qualia, the classic domain of entities from Montague Grammar can be organized as a hierarchical system of subtypes (sorts), structured into three broadly defined types:

(67) *Natural types*: Natural kind concepts grounded in the Formal and Constitutive qualia roles.
Artifactual types: Concepts grounded in the Telic (purpose or function), or Agentive (origin) qualia roles.
Complex types: Concepts integrating reference to the relation between at least two types from the other levels.

Qualia structure has proved to be an expressive representational device and has been adopted by adherents of many other grammatical frameworks. For example, Jensen and Vikner (1994) and Partee and Borschev (1998) as well as Borschev and Partee (2001) appeal to qualia structure in the interpretation of the genitive relation in NPs, while many working on the interpretation of noun compounds have developed qualia-based strategies for interpretation of noun–noun relations (Busa and Johnston, 1996; Johnston and Busa, 1999; Jackendoff, 2010). Van Valin Jr. and La Polla (1997) and Van Valin Jr. (2005) have also adopted qualia roles within several aspects of analyses in Role and Reference Grammar (RRG), where nominal semantics have required finer-grained representations.

2.8 Type theory and the lexicon

How can lexical semantic information be exploited in the compositional operations responsible for building larger semantic expressions in language? In this section we discuss how the representational strategies and approaches outlined above can be integrated into a uniform framework, one that can interface with the mechanisms we are familiar with within compositional semantics. Specifically, we present a type-theoretic treatment of lexical semantic information, as developed in Pustejovsky (2006), Asher and Pustejovsky (2006), and subsequently in Asher (2011).

The basic semantic knowledge associated with a lexical item can be split into two categories: its *semantic type* and its *selectional type*. While the former identifies the semantic class that a lexical item belongs to (such as *entities*, *events*, and *attributes*), the latter class specifies the semantic characteristics associated with the arguments to a lexical item. From a type-theoretic position, such as that adopted within formal semantics since Montague, these are in fact identical, as the type of a functional expression is determined by the nature of the typing of its arguments (see Partee, 1975, 1978; Bach, 1984).

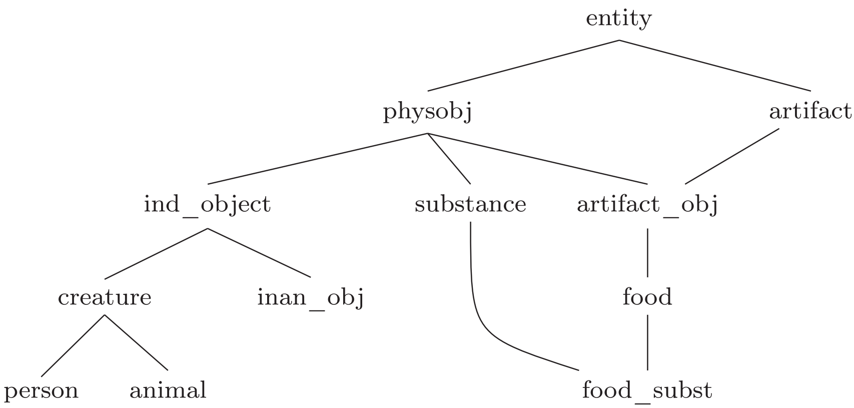

Figure 2.1 *Fragment of a type hierarchy*

Following standard assumptions within type-theoretic semantics (see Chierchia and McConnell-Ginet, 2000), we assume a vocabulary of atomic and derived types. For the present discussion, we let the atomic types be: e the general type of entities; and t the type of truth values. Derived types consist of all *functional types*, formed with the functional infix, $\rightarrow$; hence, for any types σ and τ, then $\sigma \rightarrow \tau$ is a type.

As in Montague Grammar and other standard frameworks, we will take a lexical entry to consist of a lambda term and a type assignment to the variables in the term. This will then determine via the standard interpretation for the lambda term a functional type for the whole expression. Unlike Montague Grammar, however, *Type Composition Logic* (TCL) (see Asher and Pustejovsky, 2006; Asher, 2011) has a much richer system of types reflecting the information conventionally associated with a word in the Generative Lexicon approach (its qualia structure), and correspondingly more complex rules for manipulating these types. Such typing knowledge is typically structured by means of some sort of inheritance mechanism (Evans and Gazdar, 1989, 1996; Copestake, 1993; Pollard and Sag, 1994; Copestake and Briscoe, 1997), and modeled as a lattice or semi-lattice (Carpenter, 1992). Briscoe et al. (1993) describe a rich system of types for allowing default mechanisms into lexical type descriptions. An example of such an inheritance structure is shown in Figure 2.1.

As presented in the previous section, the qualia structure is a system of relations that characterizes the semantics of a lexical item, and TCL provides a way to encode this directly in the type associated with a lexeme. Following Asher and Pustejovsky (2006), we adopt two additional type constructors, $\otimes$ and $\bullet$. The constructor $\otimes$ introduces a qualia relation, Q, such that if σ and τ are types, then so is $\sigma \otimes_Q \tau$, for Q. The constructor $\bullet$ introduces a *complex type*, such that if σ and τ are types, then so is $\sigma \bullet \tau$.

To illustrate how a richer lexical semantics is reflected directly in the type structure, notice that a feature structure denoting the qualia structure for an expression, α, reduces to the type $\alpha : (\beta \otimes_A \sigma) \otimes_T \tau$ in TCL (where A and T stand for Agentive and Telic qualia, respectively).

(68)
$$\begin{bmatrix} \alpha & \\ \text{QS} = & \begin{bmatrix} \text{FORMAL} & = \beta \\ \text{TELIC} & = \tau \\ \text{AGENTIVE} & = \sigma \end{bmatrix} \end{bmatrix}$$

We can define all three basic types from (67) using the qualia and these constructors. The first two classes in (67) are defined in terms of qualia; a natural physical object is a subtype formed from predication with the atomic type, e. The natural types, $\mathcal{N}$, are those subtypes that are formally structured as a join semi-lattice, $\langle \mathcal{N}, \sqsubseteq \rangle$, in this domain (Pustejovsky, 2001). We say that this type occupies the value of the FORMAL quale in the type structure. The creation of a predicate that, in turn, selects for a natural type follows conventional functional typing assumptions: for any type τ in the subdomain of natural types, $\tau \in \mathcal{N}$, $\tau \to t$ will be considered a *natural* functional type.

Once we have defined natural type entities, their corresponding functional types are defined. The creation of predicates over the subdomain of natural types follows conventional functional typing assumptions: for any type τ in the subdomain of natural types, $\tau \in \mathcal{N}$, $\tau \to t$ is a *natural* functional type. This allows us to define natural predicates such as *die* or *touch*, as: $\lambda x{:}e_N[die(x)]$, $\lambda y{:}e_N \lambda x{:}\ e_N[touch(x, y)]$.

The second class of types, *artifactual types*, are defined as any type with an associated TELIC type. For example, the concept of a potable liquid would have the typing, $liquid \otimes_T drink$. Similarly, an artifactual entity such as *bread* would be typed as, $(phys \otimes_A bake) \otimes_T eat$. The creation of functional types over the domain of artifactual types is defined as follows: for any type τ in the domain of artifactual entity types, $\tau \in \mathcal{A}$, $\tau \to t$ is an *artifactual* functional type. For example, $\lambda x{:}e_A[break(x)]$; $\lambda y{:}e_A \lambda x{:}e_N[fix(x, y)]$.

Finally, *complex types* are constructed through a type-construction operation (the dot, $\bullet$) over the domain of Naturals, Artifactuals, and Complex Types. Consider the noun *book*, a complex type denoting both the informational content and the physical manifestation of that content: $phys \bullet info$. Other examples include the nouns *lunch*, *school*, and *promise*. Constructing functional types over the subdomain of complex types is straightforward: for any type τ in the domain of complex entity types, $\tau \in \mathcal{C}$, $\tau \to t$ is a *complex* functional type. Examples include verbs such as *read*: $\lambda y{:}phys \bullet info$ $\lambda x{:}e_N[read(x, y)]$.

One of the advantages of introducing a finer-grained system of types is the ability to explain the selectional constraints associated with the arguments to a predicate; that is, *die* selects for an animate entity; *assemble* selects for a semantically plural individual, and so on.

Another advantage of a richer type system, however, is the way that it facilitates new solutions to semantic analyses that were not previously available. For example, consider the problem of verbal polysemy, mentioned in Section 2.3.1 above, where a verb is able to appear in multiple syntactic contexts, with complementary or additional arguments. The examples below illustrate this phenomenon with the verb *enjoy*, an aspectual verb *begin*, and an experiencer causative, *wake up*.

(69) a. Mary enjoyed *the movie*.
 b. Mary enjoyed *watching the movie*.

(70) a. Mary began *a book*.
 b. Mary began *to read a book*.

(71) a. *The coffee* woke John up.
 b. *John's drinking the coffee* woke him up.

What these sentence pairs illustrate is the process of *type coercion*, a semantic operation that converts an argument to the type which is expected by a predicate, where it would otherwise not be accepted (see Pustejovsky, 1995). This is an operation in the grammar ensuring that the selectional requirements on an argument to a predicate are satisfied by the argument in the compositional process. The rules of coercion presuppose a typed language such as that outlined above. By allowing lexical items to coerce their arguments, we obviate the enumeration of multiple entries for different senses of a word.

The notion that a predicate can specify a particular target type for its argument is a very useful one and intuitively explains the different syntactic argument forms for the verbs above. In sentences (69) and (70), noun phrases and verb phrases appear in the same argument position, somehow satisfying the type required by the verbs *enjoy* and *begin*. Similarly, in sentences (71), noun phrases of very different semantic classes appear as subject of the verb *wake*.

If we analyze the different syntactic occurrences of the above verbs as separate lexical entries, following the sense enumeration theory outlined in previous sections, we are unable to capture the underlying relatedness between these entries; namely, that no matter what the syntactic form of their arguments, the verbs seem to be interpreting all the phrases as events of some sort. It is exactly this type of complement selection which type coercion allows in the compositional process.

2.9 Open issues in lexical semantics

Contemporary views on what constitutes lexical knowledge still vary significantly in formal semantics, artificial intelligence, and computational linguistics. To a large extent, the nature of grammar itself is determined by what information the lexicon supplies to other linguistic components. Even with this accepted view, the lexicon has been viewed as a mostly passive module of grammar. What has changed over the past twenty years is the amount of semantic information that is associated with lexical items, as well as what is taken to be a lexical item in the first place. What started as a static repository of pure stems, with active word formation rules generating everything from inflected forms to derivational compounds, has developed into a rich library of lexically based syntactic encodings, generative operations, and type-theoretic processes.

One of the most challenging questions confronting lexical semantics researchers is perhaps how to integrate to and interface with discourse semantics and contextual information. If lexical meaning is context-sensitive at the sentence level, as suggested in previous sections, then perhaps lexical semantics should directly encode how words are ambiguous or underspecified at the level of discourse as well. It has been proposed (Lascarides and Copestake, 1999; Asher and Lascarides, 2003b; Bos, 2004) that context-sensitive approaches to both sentence composition and discourse interpretation should have a common view about meaning, some of the same formal tools, and some of the same problems. There is not only context sensitivity in both lexical and discourse meaning but also interdependence.

It has alternatively been suggested that this issue reduces to the question of how "pragmatic enrichment" is formally handled in the interpretation of a contextually situated utterance. Some researchers argue that pragmatic or situational frames encode "unarticulated constituents" that fill in the gaps in the interpretation of an utterance left by a syntactically impoverished sentence (Travis, 1981; Bach, 1994; Sperber and Wilson, 1995; Carston, 2002; Recanati, 2004). Countering this view, Stanley (2002a,b), Stanley and Szabó (2000), and others claim that any truth-conditional effect of extra-linguistic context to a sentence can be traced to the logical form of the lexical items involved. Both theories have data to support their arguments and there are merits to both positions. Some resolution to this question may in fact come from computational models of discourse and how lexical knowledge is exploited in context.

Acknowledgment

I would like to thank the editors and the anonymous reviewer for valuable comments and feedback on this chapter.

3

Sentential semantics

Peter Pagin

3.1 Sentences and sentence meaning

There are three basic conceptions of a *sentence*: syntactic, semantic and pragmatic (Stainton, 2000). According to the syntactic conception, a sentence is an expression with certain grammatical properties, as specified in a grammar. According to the semantic conception, a sentence is an expression with a certain type of meaning, for instance a sentence expressing a *proposition*, something that is true or false (with respect to the actual world). According to the pragmatic conception, a sentence is an expression with a certain kind of use, typically that of making a *speech act*.

These three conceptions are naturally enough pretty well correlated. Speakers of natural languages typically use sentences in the grammatical sense for making speech acts and expressing propositional thoughts by means of the sentence meaning. Nevertheless, in many cases they come apart. On the one hand, speakers often use sub-sentential expressions, such as 'Reserved for tonight', pointing to a chair (Stainton, 2000, p. 446), for making a speech act.

On the other hand, very often, what is a grammatical sentence does not have a meaning that is simply a propositional content in an ordinary sense. This can happen for a variety of reasons, such as indexicality, presupposition, conventional implicature, discourse phenomena, interrogative mood.

In this chapter, we shall be concerned with sentences in the syntactic sense, and we shall look at how semantic theories of various types model sentence meaning. In some cases we will also consider their philosophical motivations. The topic will be delimited in certain ways. We shall only discuss *declarative* sentences (see Dekker et al., Chapter 19, and Portner, Chapter 20 for chapters on non-declaratives). We shall also not cover dynamic phenomena in discourse semantics (see Asher, Chapter 4). We are also not going to discuss presupposition and similar phenomena in the semantic/pragmatics interface (see Schlenker, Chapter 22). We *shall* be concerned with semantic context dependence and related phenomena.

One of the key features of the syntactic conception of a sentence is that sentences are syntactically (or morphosyntactically) *complex*. Since they result from combining linguistic elements, there is a question of how the meaning of the sentence is related to the meanings of its *parts* and the way they are combined. This will be the topic of Section 3.3. The rest of the chapter is organized as follows: Section 3.2 gives a conceptual background for talk about meanings and semantic values. Section 3.3 presents the idea of semantic compositionality, informally and formally. Section 3.4 presents Frege's views on sentence meaning, and Section 3.5 does the same for Davidson's. Section 3.6 gives a presentation of possible-worlds semantics and is followed by sections on Montague semantics (3.7), Kaplan style semantics (Section 3.8), centred-worlds propositions (Section 3.9), relativism (Section 3.10), two-dimensionalism (Section 3.11), and situation semantics (Section 3.12). Section 3.13 is concerned with structured meanings, Section 3.14 with verificationism, and Section 3.15 with procedural semantics. Section 3.16 contains some concluding words.

3.2 Semantic values and meta-semantics

There are two main formats that a semantic theory takes with respect to sentences. First it can ascribe one or more *properties* to sentences and specify the *conditions* under which a sentence has this property or properties. Typically the property is that of *being true*, or some property similar to or related to that. Truth definitions are of this kind. Let us call this the *axiomatic* style.

The alternative is to assign explicit *semantic values* to sentences. The semantic theory then takes the format of specifying one or more *semantic functions* that map expressions on semantic values. Let us call this the *algebraic* style.

These two formats are typically inter-translatable. For instance, in a *truth-theoretic* (axiomatic) format, the meaning specification for sentences takes the form of material biconditionals of the form '*s* is true iff *p*' (see Section 3.5). Whether the sentence *s* is true, then, depends on whether *p* is the case, which is determined by the (actual) world. In a corresponding algebraic framework, we have a function μ that assigns values 1 (for truth) or 0 (for falsity) to sentences, and we can equally ask for the conditions under which the value is 1.

By contrast, whether we include possible worlds, and therefore have an *intensional* semantics, or leave possible worlds and similar entities out of the semantics makes a crucial difference. This difference, in turn, is related to the distinction between semantics and meta-semantics. In an *extensional* semantics, the standard semantic value of a one-place predicate is a set of objects – informally, the set of objects of which the predicate is true.

The conditions for being in the extension of the one and for being in the extension of the other, however, are intuitively distinct. You can take account of this in two different ways. The first is semantic. You add intensions to your semantic theory, typically as functions from possible worlds to extensions. Then you have a theory with two levels: on one (upper) level intensions, and on the other (lower) level extensions with respect to worlds.

The second way is meta-semantic. Semantics proper only deals with extensions, but there are certain facts in the world which combine to determine what the semantic value is. These facts, how they combine and how they select the semantic value, on this view fall strictly outside semantics. The debate concerning the meaning of proper names, between Fregeans and Millians, exhibits this conflict over meaning determination. For Fregeans it belongs to semantics (*sense*), while for Millians it does not.

Many semantic frameworks include more than one type of semantic value or property, often but not always at different levels. Whether the framework has certain formal properties, such as compositionality, can then be asked for each of the types of value separately (is *extension* compositional; is *intension* compositional?). In some cases the different kinds of value interact in such a way that it is rather the system as a whole that must be characterized, as we shall see.

3.3 Compositionality

There is a remarkable fact about human linguistic communication: language users manage to successfully communicate new thought contents (contents not thought before) by means of new sentences (not used or considered before). Since mind-reading is out of the question, and *chance* would deliver a microscopic rate of communicative success (given how many sentences and contents there are), there must be something about language itself that makes this possible, as was also noted by Frege (1914, 1923):

> It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a terrestrial being for the very first time can be put into a form of words which will be understood by somebody to whom the thought is entirely new. This would be impossible, were we not able to distinguish parts in the thought corresponding to the parts of a sentence, so that the structure of the sentence serves as an image of the structure of the thought. (Frege, 1923, opening paragraph)

The hearer's task in communication is that of finding an appropriate *interpretation* of an utterance. On a widely accepted view, the ability to discharge this task for *new* sentences is made possible by a systematic correlation between the syntactic build-up of sentences and their meanings. It is common to spell out this correlation by means of the principle of *compositionality* (PC):

(PC) The meaning of a complex expression is a function of the meanings of its parts and its mode of composition.
(First stated in this format by Partee, 1984a, p. 281)

We shall say that the *semantics* of language is compositional or not, intending the assignment of meaning by some particular semantic function or semantic theory. In a derivative sense, we can say that the *meaning property* of a particular language is compositional if the *correct* semantics for it is compositional in the first sense.

In interpreting (1a), given the syntactic structure of (1b),

(1) a. John saw the dog
b. $[_{S}\ \text{John}_{NP}\ [_{VP}\ \text{saw}_{VT}\ [_{NP}\ \text{the}_{Det}\ \text{dog}_{N}]]]$

the hearer knows the meaning of *the*, *dog*, *John*, and *saw*, and the semantic significance of (a) the Determiner + Noun (Det+N) combination, (b) the Verb (transitive) + Noun Phrase combination (VT+NP), and (c) the Noun Phrase + Verb Phrase combination (NP+VP). He starts with working out the meaning of *the dog*, using (a), goes on working out the meaning of *saw the dog*, using (b), and finishes by working out the meaning of (1a) from there, using (c).

Although (PC) is a standard formulation, what is meant by 'part', as the principle is normally understood, is *immediate part* (immediate constituent). The principle is often incorrectly stated as saying that the meaning of the complex expression is a function of the meanings of the *simple parts* and the way they are combined. But this principle is degenerate. The mode of combination of the *simple* parts in (1) is in itself a combination of Det+N, VT+NP, and NP+VP constructions, and the speaker would have to know the semantic significance of that complex mode. In fact, there are denumerably many ways of combining simple parts, and so if the hearer needs to know the semantic significance of each, the explanation of how he can understand new sentences cannot any more appeal to a working out strategy based on *finite* knowledge.

A principle closely related to (PC) is the principle of *substitutivity* (PS) (where '$A[e'/e]$' designates the expression that results from substituting e' for e in one or more occurrences).

(PS) If in a complex expression A a constituent e is replaced by a constituent e' that has *the same* meaning as e, then the meaning of the new expression $A' = A[e'/e]$ has the same meaning as A.

The intersubstituted constituents need not be immediate. In fact, if the semantics is *total*, so that each grammatical expression is meaningful, then (PC) and (PS) are *equivalent*.

If, by contrast, the semantics is *partial*, the two principles can come apart. On the one hand, if not all parts of a complex expression are meaningful, then the meaning of the complex is trivially not a function of the meanings of the parts ((PC) does not hold), but it can still hold that when two parts with the same meaning (i.e. meaningful parts) are intersubstituted, the meaning of the complex is preserved ((PS) holds). On the other hand, it may also be that A is meaningful, and that e and e' mean the same, yet $A[e'/e]$ is *not* meaningful ((PS) does not hold), while it is still true that the meaning of every complex expression that *has* meaning is a function of the meaning of its parts and their mode of combination ((PC) holds).

These observations are due to Wilfrid Hodges (2001). He calls the principle that the parts of every meaningful expression are meaningful *the domain principle*. The principle that if two expressions mean the same, then substituting the one for the other does not lead to loss of meaningfulness in the complex is called *the Husserl principle* (and the semantics *husserlian*) by Hodges. He also proved that given that the domain principle and the Husserl principle hold, (PC) and (PS) are again equivalent, even if the semantics is partial.

Using the equivalence, we can see that (PC) in several respects is quite weak. To get a counterexample to (PS), and thereby to (PC), we need two expressions e and e' with the *same* meaning, and two complex expressions, A and $A[e'/e]$ with *different* meanings. If there is no counterexample, (PS) and hence also (PC) hold. Now, if no two different expressions in the language in question have the same meaning (meaning is *hyperdistinct*), then no counterexample is possible, and hence the semantics is vacuously compositional. Similarly, if no two expressions in the language differ in meaning, again the semantics is trivially compositional.

The first of these two observations, that hyperdistinctness entails compositionality, is somewhat damaging for the explanatory power of compositionality. It may be that the meaning of complex expressions of a language obeys no regularity whatsoever, so that it is impossible to *work out* the meaning of a new complex expression, and yet because of hyperdistinctness, meaning is compositional. So compositionality alone does not explain how speakers can work out the meaning of new complex expressions. Some further property related to processing is required (see Pagin, 2012).

So far we have been talking as if semantic values or properties are assigned to *expressions*, where expressions are understood as *surface* forms, that is, types of spoken utterance, types of written inscriptions, or types of other concrete marks or events. Surface forms are, however, often syntactically ambiguous. Two different words may have the same surface form (homonymy), and complex expressions may on top be structurally ambiguous. Since different disambiguations often yield different meanings, we must define the semantic function for *disambiguated expressions*, such as analysis trees, or *grammatical terms*.

For setting out compositionality formally, we therefore need a syntactic framework, which specifies disambiguated expressions of the language, together with a domain of possible semantic values. Then the semantics is defined as a function from disambiguated expressions to semantic values. In the modern tradition, there are two approaches to this. Both are algebraic, in the sense that the syntax is conceived of as an *algebra* with syntactic operations defined over a basic set of expressions. In the tradition from Montague (1970b, 1973b), Janssen (1986, 1997) and Hendriks (2001), the syntax is built up from *categories* of expressions and corresponding sets of expressions of these categories. There are also rules that specify the categories of the arguments and values of syntactic operations. We shall look more closely at this idea in Section 3.7 on Montague semantics.

In the more recent tradition from Hodges (2001), there is no appeal to categories and types. Instead, syntactic operations are taken to be *partial*; they are defined only for certain arguments (combinations of arguments). We shall here follow Hodges, with some modifications. The syntax for a language L is a triple (G_L, E_L, V_L), where G_L is a *grammatical term algebra*, E_L is the set of *expressions* of L, and V_L is a mapping from grammatical terms to expressions (for convenience we shall henceforth drop the subscript). G itself is a triple (T, Σ, A), where T is the set of *grammatical terms*, Σ the (finite) set of *operations* that map n-tuples of grammatical terms on grammatical terms, and A is a (finite) set of *atomic* grammatical terms. T is the *closure* of A under Σ, that is the set of terms that are generated from A by means of (possibly repeated) applications of the operations in Σ.

To exemplify, let A be the set {'John', 'saw', 'the', 'dog'} and Σ the set $\{\alpha, \beta, \gamma\}$, corresponding to the rules for forming Det+NP, VT+NP, and NP+VP, respectively, of Example (1). Then the grammatical term that corresponds to the phrase structure analysis (1b) is as follows:

(2) γ('John', β('saw', α('the', 'dog')))

The V function maps the grammatical terms on the expressions of L. For each operation σ on terms, there is a corresponding operation $\overline{\sigma}$ on expressions such that

(V) $V(\sigma(t_1, \ldots, t_m)) = \overline{\sigma}(V(t_1), \ldots, V(t_n))$

In the very simple case of our example, we just have the equation in (3) where σ is α or β or γ, and '␣' marks a word boundary (space).

(3) $V(\sigma(t_1, t_2)) = V(t_1)␣V(t_2)$

After three applications of (V), we therefore get the following in (4):

(4) $V(\gamma($'John', $\beta($'saw', $\alpha($'the', 'dog'$))))$ = 'John'␣'saw'␣'the'␣'dog'

This is a structural-descriptive name of *John saw the dog*.

No abstract limitations are imposed on V (such as that the value of the first argument should appear in the expression (surface string) to the left of the value of the second argument), although a number of restrictions will be empirically motivated, in order that the hearer be able to parse the string. The algebraic framework as such is very general; it can be adapted to phrase structure rules as well as to transformations or movements.

A *semantic* function μ is also defined on the grammatical terms, mapping terms on a domain M of semantic values. Given the algebraic framework, we can now state a more formal version of (PC):

(PC′) For every n-ary syntactic operation $\sigma \in \Sigma$, there is a function $r_\sigma : M^n \longrightarrow M$ such that for all grammatical terms $t_1, \ldots, t_n$ such that $\alpha(t_1, \ldots, t_n)$ is defined and μ-meaningful, it holds that $\mu(\alpha(t_1, \ldots, t_n)) = r_\sigma(\mu(t_1), \ldots, \mu(t_n))$.

We call the function r_σ a *meaning operation*. If (PC′) holds for a semantic function μ, we say that μ is compositional. Compositionality is a general formal property that any particular semantic function, *given* a syntax, either has or lacks.

Historically, the first precise statement of (PS) was given by Frege (1892b, p. 32), for *Bedeutung*. Clarity about (PC) developed slowly after that and did not reach maturity until the mid-to-late 1970s. In the following sections we shall see how this definition applies to various types of semantic theory. In some cases, we will find reason to go beyond the standard compositionality expressed in PC′. For more on compositionality, see Janssen (1997), Szabó (2008), Pagin and Westerståhl (2010a,b, 2011).

3.4 Frege

During the seventeenth, eighteenth, and nineteenth centuries, the idea of *judging* something to be the case and *what* one judges to be the case were not sharply distinguished. On the contrary, it was assumed that a judgment has at least two components (a *subject idea* and a *predicate idea*), and it is precisely through the *judging*, and in particular the *affirming*, that the two ideas are combined into a complex whole.

Anticipated by Bolzano, Frege insisted that the judging activity and content must be separated. He argued in the late essay ‘Negation’ (Frege, 1918b) that logically complex sentences, such as negations and conditionals, when affirmed, contain subsentences with contents that may be denied (in the case of negation) or not judged at all (in the case of the conditional). Since the content of the complex sentence contains parts that correspond to the subsentences, the subsentences too must have contents, even though they are not affirmed. Having content must therefore be distinct from being affirmed.

The view itself was already in place in Frege’s middle period (especially Frege, 1892b), when he developed his mature ideas of linguistic meaning. Frege employed two semantic values, *reference* (*Bedeutung*) and *sense* (*Sinn*). These notions are introduced at the beginning of Frege (1892b), applied to the case of proper names, in order to solve the problem of the difference between informative and uninformative true identity statements. The concept of the sense of an expression is explained to be that which contains the *mode of presentation* of the referent.

In the case of a proper name, the reference is the bearer of the name. To exemplify different modes of presentation of the bearer, Frege uses definite descriptions. For instance, for the name *Aristotle*, Frege suggests ‘the student of Plato and teacher of Alexander the Great’ (Frege, 1892b, note 2) as articulating the sense/mode of presentation. However, Frege does not say that it is invariably the case that the name is synonymous with some (non-trivially formed) definite description.

For Frege, truth values, The True and The False, are objects. The sense of a sentence, then, what Frege called a *Thought* (ein *Gedanke*), is then (what contains) a mode of presentation of either The True or The False. So as in most theories of sentence meaning, the non-extensional semantic value is defined in relation to truth and falsity. However, Frege’s theory is not truth-conditional in any ordinary sense. *Sense* is not meant to encode the conditions under which the sentence refers to The True.

Frege does not specify the relation between sense and reference in purely semantic or metaphysical terms, but rather in *cognitive* terms. For instance, as spelled out by Perry (1977), Frege appears to accept the following as a *criterion of difference of sense*:

(Dif) If A understands S and S', and accepts S as true while not accepting S', then S and S' have different senses.

So, sense is as fine-grained as the discriminations that subjects make in their attitudes. This combines well with his semantic view about difference of sense. As mentioned above, Frege stated the substitution principle for Bedeutung (see (PS), p. 69), but did not state it separately for Sinn. He did, however, seem to believe in the substitution version of *inverse compositionality* for sense:

(PIC) For any complex expressions A and A', if A and A' have the same meaning, then A' results from A through replacing some sub-expressions $e_1, \ldots, e_n$ by some expressions $e'_1, \ldots, e'_n$, respectively, where e'_i has the same meaning as e_i, for $1 \leq i \leq 1$.

(PIC) is a sufficient condition for difference of meaning. Given two complex expressions A and A', if their top structure, i.e., the mode of composition of their immediate constituents, are different, then they differ in meaning, and if any of their corresponding immediate constituents differ in meaning, then A and A' differ in meaning as well. As pointed out by Heck and May (2011), Frege seems to have believed in the inverse compositionality for *sense*. For Frege it was important that mathematical expressions differing as much as e.g. '4×4' and '2^4' also differ in sense (discussed especially in Frege, 1891). Compositionality does not force the expressions to differ in sense, only allows them to.

The (Dif) principle is closely related to Frege's semantics for attitude contexts, or more generally, what he called *indirect contexts*. Indirect contexts include *oratio obliqua* contexts and attitude contexts. It is part of Frege's doctrines about sense (1892b) that expressions occurring in indirect contexts have an *indirect reference* and that the indirect reference is the same as the ordinary sense. Thus, in (5) the embedded sentence *Hesperus is a planet* does not have its ordinary reference (The True), but its indirect reference, which is the Thought *that* Hesperus is a planet.

(5) John believes that Hesperus is a planet.

This doctrine of indirect reference avoids violations of the (PS) principle for reference. For it may be that (5) is true, while (6) is false, despite the coreference of *Hesperus* and *Phosphorus*:

(6) John believes that Phosphorus is a planet.

This would violate the (PS) principle for reference, since we get (6) from (5) by substituting *Phosphorus* for *Hesperus*. By the doctrine of indirect reference, the two names differ in reference in this context, given that they already differ in sense.

Nonetheless, the standard principles (PC) and (PS) of compositionality are not satisfied for reference in Frege's semantics, since they do not make provision for the idea that the semantic value of an occurrence of an expression may depend on the *linguistic context* of that occurrence. In fact, Frege's semantics makes the reference of belief sentences depend on the *senses* of some of its part. Letting μ_r be the reference function and μ_s the sense function, we can state the general clause for belief sentences as follows:

(7) $\mu_r(B(N, p)) = r_b(\mu_r(N), \mu_s(p))$

Here '$B(N, p)$' is short for 'N believes p', and r_b is a function such that $r_b(x, y) =$ True iff x is belief-related to y (where y is a Thought).

This clause is clearly not compositional, since there is a *switch* from μ_r to μ_s for the second argument to B. There is, however, a generalization of standard compositionality that accepts such switches. *General compositionality* holds for a set $S = \{\mu_1, \ldots, \mu_n\}$ of semantic functions, where one element, say μ_1, is *designated* as applying to *unembedded* occurrences of expression. So semantic evaluation starts with μ_1. The semantics makes use of a *selection function* $\Psi : S \times \Sigma \times N \longrightarrow S$ that maps a semantic function, a syntactic operation, and an argument place on a semantic function. The principle of general compositionality can then be stated as follows:

(GPC) A system $S = \{\mu_1, \ldots, \mu_n\}$ of semantic functions is *general compositional* iff there is a selection function Ψ such that for every k-ary operation α and every member $\mu_i \in S$ there is a meaning operation r_{α,μ_i} such that for any terms $t_1, \ldots, t_k$, if $\alpha(t_1, \ldots, t_k)$ is μ_i-meaningful, then

$$\mu_i(\alpha(t_1, \ldots, t_k)) = r_{(\alpha,\mu_i)}(\mu_{j_1}(t_1), \ldots, \mu_{j_k}(t_k))$$

where $\mu_{j_m} = \Psi(\mu_i, \alpha, m)$, $1 \leq m \leq k$.

Standard compositionality is the special case where S has just one element (cf. Pagin and Westerståhl, 2010c).

As stated in (GPC), a general compositional system S has finitely many members. There is a question whether Frege's semantics meets this condition. In the equation (7), the semantics only makes use of the fact that the reference of the belief sentence depends on the sense of the embedded sentence, not on the additional idea that the sense of the embedded sentence also is the reference, *in that context*. Given Frege's idea that the sense of an expression *presents the referent*, and hence *determines the referent*, it may seem that the sense of the *embedded* occurrence of the embedded sentence in (5), *Hesperus is a planet*, must have a sense different from its ordinary sense, an *indirect sense*. It is therefore a common interpretation of Frege that there is an infinite *hierarchy* of indirect senses. In a simply iterated belief sentence like (8), the sentence (5) occurs in an indirect context, and so has indirect reference.

(8) Bill believes that John believes that Hesperus is a planet.

The sentence *Hesperus is a planet* occurs *doubly embedded* and so occurs in an indirect indirect context. On some interpretations of Frege, the sentence then has a doubly indirect reference. The doubly indirect reference, in turn, is the *simply indirect sense* it has as embedded in (5). Doubly embedded in (8), it has a doubly indirect sense. And so on. Using superscript indices for level, where $\mu^0 = \mu_r$ and $\mu^1 = \mu_s$, and a corresponding index on the meaning operations r^i_b, we would have the equation in (9):

(9) $\mu^i(B(N, p)) = r^i_b(\mu^i(N), \mu^{i+1}(p))$

There is some textual evidence that Frege actually believed in indirect senses (see Frege, 1902). With infinitely many senses μ^j, $j = 0, 1, 2, \ldots$, the condition of general compositionality that S is a finite set is not met. Frege has also been criticized in the literature (e.g. in Davidson, 1965) for providing an *unlearnable semantics*. Note that it is not the semantic switching in itself that motivates the hierarchy, but the ideas that the sense plays the double roles of reference in the indirect context *and* reference determination.

One last consequence of this idea deserves to be mentioned. In Frege's semantics, syntactic combination invariably has the form of letting one expression fill an *argument place* in another expression. A one-place predicate is an *unsaturated expression*, that is, it has an argument place, and in this case for a singular term. When the argument is given, we get a sentence. Correspondingly, the semantic significance of such a combination is *function application*: the reference of the unsaturated expression is a function, which applies to the reference of the argument expression. Thus, we have in general the following (AP = application):

(AP) $\mu_r(\alpha(t)) = \mu_r(\alpha)(\mu_r(t))$

Now, although Frege did not claim that sense is compositional, the idea that reference is compositional (modulo context) and that ordinary sense is a kind of reference naturally leads to the claim that sense is compositional (modulo context) as well. That is, we should, at least in all cases that do not involve context shift, apply (AP) to sense as well (just substitute 's' for 'r').

However, this idea is problematic. It cannot plausibly be combined with the idea, which Frege also had (e.g. in Frege, 1923), that the sense of a *part* of an expression is a *part* of the sense of the whole. For if the sense of α is a function that applied to the sense of t gives as value the sense of $\alpha(t)$, which in turn contains the sense of α as a proper part, then the sense of α is a non-well-founded entity: it is a function defined partly in terms of itself (see Dummett, 1981a, p. 293).

All in all, Frege's conception of the sense of a sentence contains several strands that are in conflict with each other. This comes out even more in his account of indexical expressions (Frege, 1918a).

3.5 Davidsonian truth-theoretic semantics

On the face of it, Davidson-style truth-theoretic semantics is a straightforward truth-conditional semantics, where sentence meaning is specified in the following format:

s is true in L iff p.

However, things are more complicated. On one level, there is a semantic theory of an axiomatic kind, which uses only extensional concepts. On a second level, there are meta-semantic ideas about the *modal force* of the theorems of such a theory. On a third level, again also meta-semantic, there is an idea that meaning is captured by *sets* of such theories. On top, there is an elaborate meta-semantic theory, of *radical interpretation*, for selecting the right set of theories. This presentation will be based on the papers collected in Davidson (1984).

A truth-theory T is intended to have the form of a Tarskian truth definition but is meant to be an empirical theory about a natural language *L*. The language is identified as the language spoken by a particular speaker, (call her S). Davidson (1965, 1967b) argues that a meaning theory must show how it is possible for a speaker with finite knowledge to have infinite capacity, that is, the capacity to understand in principle each of the infinitely many meaningful sentences of a natural language. Although he does not use the term, his general requirements are, nearly enough, that such a theory should be compositional. More specifically, it should be an *axiomatic* theory with the following components:

(i) Axioms for simple singular terms
(ii) Axioms for simple predicates
(iii) Recursive semantic clauses for constructions
(iv) Theorems for sentences

For each sentence *s* of the object language, the theory must contain a theorem that specifies the relevant semantic property of *s*. The theorems are to be derivable from the axioms, by means of the recursive clauses, and a suitable set of rules of inference.

He also argues (Davidson, 1967b, pp. 20–22) that such a theory should use only extensional semantic vocabulary. The reason is that intensional vocabulary, such as 'means that', creates non-extensional contexts that create problems for the derivations, rendering some needed rules of inference invalid. He suggests that the proper way of satisfying these requirements is to use the Tarskian format of a truth definition. The semantic ingredients are, then, as follows:

(i) a *reference* function *R* that assigns referents to simple closed terms;
(ii) a *satisfaction* relation SAT relating infinite sequences of objects and open sentences (possibly containing free variables);
(iii) a truth property.

All of these are relativized to an object language. Davidson's official position here is to assume a prior understanding of *truth*, while *reference* and *satisfaction* are *instrumental* notions, used only for deriving the theorems. The set of theorems, specifying sentence meanings, makes up the entire empirical content of such a semantic theory, called a *T-theory*. For the simple case of the language of predicate logic, we set out a version of the format of the clauses below. The symbol s with subscripts ranges over open and closed sentences; σ with subscripts ranges over infinite sequences of objects, and $\sigma_i(j)$ is the object in the j:th position of sequence σ_i. We use $\sigma[a/b]$ to denote the sequence that results from σ by replacing object b, wherever it occurs in σ, by object a. We use x with subscripts for object language variables, and t with subscripts for object language closed singular terms. F_i is used for an object language predicate, and G_i for the corresponding meta-language predicate. $A(t_1, \ldots, t_n)$ is an object language (open or closed) sentence that contains the terms $t_1, \ldots, t_n$. We use $\vdash_T$ to indicate that the sentence that follows is derivable in the T-theory. Then we have the following:

(TT) Clauses in a T-theory for a first-order language:
- (i) $\vdash_T R(t) = a$
- (ii) $\vdash_T \textsc{sat}(\sigma, F_i(x_1, \ldots, x_n))$ iff $G_i(\sigma(1), \ldots, \sigma(n))$
- (iii) $\vdash_T \textsc{sat}(\sigma, At_1, \ldots, t_n)$ iff $\textsc{sat}(\sigma[R(t_i)/\sigma_i], Ax_1, \ldots, x_n)$, $1 \leq i \leq n$
- (iv) $\vdash_T \textsc{sat}(\sigma, s_1 \,\&\, s_2)$ iff $\textsc{sat}(\sigma, s_1)$ and $\textsc{sat}(\sigma, s_2)$
- (v) $\vdash_T \textsc{sat}(\sigma, \neg s)$ iff it is not the case that $\textsc{sat}(\sigma, s)$
- (vi) $\vdash_T \textsc{sat}(\sigma, \exists x_i A)$ iff there is an object a such that $\textsc{sat}(\sigma[a/\sigma(i)], A)$
- (vii) $\vdash_T \mathrm{T}(s)$ iff for all sequences σ, $\textsc{sat}(\sigma, s)$.

(i) gives the general form of reference axioms for simple terms, and (ii) gives the form of satisfaction axioms for simple predicates. (iii) provides the semantics for predication, that is, atomic sentences. (iv) and (v) give the recursive clauses for conjunction and negation, while (vi) gives the clause for the existential quantifier. Finally, (vii) equates truth with satisfaction by all sequences, just as in Tarski. We derive a theorem, a so-called T-sentence, by starting with (vii), and then applying rules according to the form of the object language sentence, substituting on the right-hand-side, until the right-hand side no longer contains any occurrence of ‘SAT’ or ‘R’.

Together with these clauses, we need to specify the inference rules that apply. It cannot be the case that the T-theory is closed under all *logically valid* rules, because that would generate many theorems that would not be of the right kind. Davidson (1973, p. 138) suggested what he called ‘canonical proofs’. He did not spell out the details, but suggested that a canonical proof ‘moves through a string of biconditionals’. In particular, a rule of substitution must be added: if it holds that $\vdash_T p$ iff q, then ‘p’ and ‘q’ may be inter-substituted in $\vdash_T$ contexts. By these means, we will be able to derive T-sentences of the form below:

$$\vdash_T \mathrm{T}(s) \text{ iff } p$$

Davidson and others set out a program to develop the truth-theoretic framework to handle other features of natural language. Davidson himself adapted it to simple context dependence, quotation, and indirect discourse. Others have extended it much further, applying it to other natural language constructions, and spelling out proof procedures (see Larson and Segal, 1995; Lepore and Ludwig, 2005).

From the point of view of capturing natural language meaning, there is a strong suspicion that T-theories, being extensional, are much too weak. Davidson himself pointed out that a sentence like (10) below is true, since both the left-hand-side and the right-hand-side of the biconditional are true, and since the biconditional is material, nothing more than that is needed.

(10) 'Grass is green' is true in English iff snow is white.

Part of the answer is to point out that it is not enough for a biconditional of the theorem form to be true in order to qualify as being a true T-*theorem*. In order to be a true theorem the T-sentence must be derivable, in the canonical way, in a T-theory *all of whose* theorems are true. That will weed out all individual T-sentences that just happen to be true. But it is not enough to weed out systematic error (see Foster's objection below on p. 79).

Another part of the answer by Davidson (1967b, 1973, 1974) consisted in the appeal to meta-semantic theory of *radical interpretation*, which selects T-theories. Radical interpretation relates T-theories to belief ascriptions (in the first place), by means of inferences of the following kind.

(11) (1) Speaker S holds sentence *s* true.
(2) $\vdash_T$ *s* is true iff *p*
(3) Hence, S believes that *p*.

For the radical interpreter, who tries to find an appropriate T-theory for the speaker S, premise (1) specifies part of the *data*; without knowing what *s* means the interpreter can notice, on the basis of the speaker's behaviour (like overt assent), that she or he holds a sentence true. Premise (2) is provided by the current semantic *hypothesis* of the interpreter. The conclusion is a belief attribution to S. If that attribution is reasonable, this confirms the semantic hypothesis. If it is *un*-reasonable, the hypothesis is *disconfirmed*.

What are reasonable belief attributions? The basic observation is that too much attribution of false beliefs and inconsistent beliefs undermines an interpretation that leads to them via inferences of type (11). *The principle of charity* says that *correct* interpretations *maximize* or *optimize* attributions of true and consistent beliefs. Charity is applied to the theory as a whole: only the best theories (those that are not inferior to others) are acceptable. The principle of charity is a criterion of acceptability, not primarily a heuristic for finding the best theories.

The fact that a language contains context-dependent sentences helps select theories by appeal to charity: sentences held true in some contexts but not in others help confirm theories that correlate them with true beliefs about the current context.

Those theories that are acceptable by appeal to charity are sometimes called *interpretive*. So, it is part of the Davidsonian concept of sentence meaning that sentence meaning is specified by a T-sentence provided it is a theorem of an interpretive T-theory.

This still leaves some loopholes, however. One was suggested by Foster (1976). Suppose we have an interpretive T-theory T_1. Then, we devise another theory T_2, so that they are related for any *s* and *p* as follows:

(12) a. $\vdash_{T_1} T(s)$ iff p

b. $\vdash_{T_2} T(s)$ iff the Earth moves and p

Given that the Earth *does* move, the two theories are equally well supported by charity, and should both be correct if one of them is. According to Davidson's reply (Davidson, 1976), acceptable T-theorems have a kind of modal force: they are *lawlike*. That they are lawlike means that they are *counterfactual supporting*: if the speaker *were to hold s* true, she or he *would* believe that *p*. The upshot is that although T_2 is true if T_1 is true, its theorems can fail to be lawlike. Thus, it is part of Davidson's conception of sentence meaning that it creates a nomic connection between holding true and belief.

Finally, since there is no guarantee that there is a unique best theory, by the charity criterion, how do we make a selection among the best? We do not. According to Davidson, they are *all* correct. They may appear on the surface to be mutually incompatible, but according to Davidson the true empirical content of interpretive T-theories is precisely what is *common* to them. The underlying facts about the speakers' attitudes are captured differently by different theories, but they are equivalent in capturing the same underlying facts. Davidson (1979) compares the situation to measuring temperature with different scales, Centigrade and Fahrenheit; different numerical values are offset by the difference in scales and so capture the same physical states. This is Davidson's *indeterminacy* of interpretation. The meaning of a sentence is precisely what is shared between what all correct (interpretive) T-theories say about it.

3.6 General possible-worlds semantics

Possible-worlds semantics is a purely semantic framework. It is not associated with any particular meta-semantic theory or any particular foundational view. The idea of possible worlds was first suggested by Leibniz. In the twentieth century it was reintroduced by Wittgenstein (1922; the space of logical possibilities) and Carnap (1956, first edition, 1947). For Carnap, a possible world was a total consistent state-description, hence a linguistic item. The possible-worlds framework was developed in the late 1950s by the addition of a so-called accessibility relation, due to Hintikka (1962), Kanger (1957c), and Kripke (1963).

The possible-worlds framework is used both descriptively and model-theoretically. In model theory, the main purpose is to give a semantics for modal operators and characterize logical properties of sentences containing them. Descriptive semantics assumes that a domain of possible worlds is given. The nature of possible worlds is largely a metaphysical question.

The basic semantic framework associates each non-logical expression *e* with semantic properties on two levels: *e* has an extension with respect to every possible world, and it has an intension, which is a *function* from possible worlds to extensions. If the semantics is partial, the intension may be undefined for some worlds; for instance, the intension of a proper name might be undefined with respect to worlds where the actual world referent of the name does not exist. Note that the intension is the *value* of the semantic function (μ), not the semantic function itself; the intension takes worlds as arguments, not expressions. Since we get extensions by application of the intension, one semantic function is sufficient.

Sentence meaning is then modelled as a function from possible worlds to truth values (sometimes as sets of possible worlds – those worlds where the sentence is true). This is the basic notion of a *possible-worlds proposition*. In a truth definition format, the clause for the necessity operator $\Box$ is that $\Box p$ is true at a world $\boldsymbol{w}$ iff p is true at all worlds $\boldsymbol{w}'$ accessible from $\boldsymbol{w}$, where R is the accessibility relation:

(Nec) $\models_{\boldsymbol{w}} \Box p \quad \text{iff} \quad \models_{\boldsymbol{w}'} p \quad$ for all worlds $\boldsymbol{w}'$ such that $\boldsymbol{w}R\boldsymbol{w}'$

Where $\mu(p)$ is an *intension* I, the corresponding intension for $\Box p$ is a function I′ such that for a world $\boldsymbol{w}$, $\mathrm{I}'(\boldsymbol{w}) = 1$ iff $\mathrm{I}(\boldsymbol{w}') = 1$ for all worlds $\boldsymbol{w}'$ accessible from $\boldsymbol{w}$. To render this in the algebraic format we define a function r_{N} as in (13), where I is an intension.

(13) $r_{\mathrm{N}}(\mathrm{I})(\boldsymbol{w}) = 1$ iff $\mathrm{I}(\boldsymbol{w}') = 1$, for all worlds $\boldsymbol{w}'$ accessible from $\boldsymbol{w}$.

And then the necessity clause is simply as in (Nec′).

(Nec′) (i) $\mu(\Box p)(\boldsymbol{w}) = r_{\mathrm{N}}(\mu(p))(\boldsymbol{w})$
(ii) $\mu(\Box p) = r_{\mathrm{N}}(\mu(p))$

Here the first sub-clause gives the application of the intension to a world $\boldsymbol{w}$, and the second sub-clause gives the equivalent non-applied form. It clearly conforms to the compositional format.

One question here is whether possible-worlds semantics satisfies the condition of giving the content of attitudes by means of the truth conditions of sentences. If Bill holds true (14a), is the content of his belief really specified by (14b)?

(14) a. John may be a secret agent.
b. There is a world $\boldsymbol{w}$ accessible from the actual world $\boldsymbol{a}$ such that John is a secret agent in $\boldsymbol{w}$,