

Blackwell Textbooks in Linguistics

Introduction to Pragmatics



Betty J. Birner

 WILEY-BLACKWELL

Introduction to Pragmatics

Blackwell Textbooks in Linguistics

The books included in this series provide comprehensive accounts of some of the most central and most rapidly developing areas of research in linguistics. Intended primarily for introductory and post-introductory students, they include exercises, discussion points, and suggestions for further reading.

1. Liliane Haegeman, *Introduction to Government and Binding Theory* (Second Edition)
2. Andrew Spencer, *Morphological Theory*
3. Helen Goodluck, *Language Acquisition*
4. Ronald Wardhaugh, *An Introduction to Sociolinguistics* (Sixth Edition)
5. Martin Atkinson, *Children's Syntax*
6. Diane Blakemore, *Understanding Utterances*
7. Michael Kenstowicz, *Phonology in Generative Grammar*
8. Deborah Schiffrin, *Approaches to Discourse*
9. John Clark, Colin Yallop, and Janet Fletcher, *An Introduction to Phonetics and Phonology* (Third Edition)
10. Natsuko Tsujimura, *An Introduction to Japanese Linguistics* (Second Edition)
11. Robert D. Borsley, *Modern Phrase Structure Grammar*
12. Nigel Fabb, *Linguistics and Literature*
13. Irene Heim and Angelika Kratzer, *Semantics in Generative Grammar*
14. Liliane Haegeman and Jacqueline Guéron, *English Grammar: A Generative Perspective*
15. Stephen Crain and Diane Lillo-Martin, *An Introduction to Linguistic Theory and Language Acquisition*
16. Joan Bresnan, *Lexical-Functional Syntax*
17. Barbara A. Fennell, *A History of English: A Sociolinguistic Approach*
18. Henry Rogers, *Writing Systems: A Linguistic Approach*
19. Benjamin W. Fortson IV, *Indo-European Language and Culture: An Introduction* (Second Edition)
20. Liliane Haegeman, *Thinking Syntactically: A Guide to Argumentation and Analysis*
21. Mark Hale, *Historical Linguistics: Theory and Method*
22. Henning Reetz and Allard Jongman, *Phonetics: Transcription, Production, Acoustics and Perception*
23. Bruce Hayes, *Introductory Phonology*
24. Betty J. Birner, *Introduction to Pragmatics*

Introduction to Pragmatics

Betty J. Birner

This edition first published 2013
© 2013 Betty J. Birner

Blackwell Publishing was acquired by John Wiley & Sons in February 2007. Blackwell's publishing program has been merged with Wiley's global Scientific, Technical, and Medical business to form Wiley-Blackwell.

Registered Office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

Editorial Offices

350 Main Street, Malden, MA 02148-5020, USA

9600 Garsington Road, Oxford, OX4 2DQ, UK

The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, UK

For details of our global editorial offices, for customer services, and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com/wiley-blackwell.

The right of Betty J. Birner to be identified as the author of this work has been asserted in accordance with the UK Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Birner, Betty J.

Introduction to pragmatics / Betty J. Birner.

pages cm

Includes bibliographical references and index.

ISBN 978-1-4051-7582-1 (cloth) – ISBN 978-1-4051-7583-8 (pbk.)

1. Pragmatics.

P99.4.P72B57 2013

401'.45–dc23

2012005347

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

Cover image: Paul Klee, *Rainy Day* (detail), 1931 (no 150), oil and pen & brush and coloured ink on gessoed burlap. Private Collection/ Photo © Christie's Images/ The Bridgeman Art Library. Cover design by Nicki Averill Design.

Set in 10/13 pt Sabon by Toppan Best-set Premedia Limited

For Andrew and Suzanne,
my two favorite people

Contents

Preface	x
Acknowledgments	xii
1 Defining Pragmatics	1
1.1 Pragmatics and Natural Language	2
1.2 The Boundary Between Semantics and Pragmatics	9
1.3 Summary	34
1.4 Exercises and Discussion Questions	36
2 Gricean Implicature	40
2.1 The Cooperative Principle	41
2.2 Types of Implicature	62
2.3 Testing for Implicature	68
2.4 The Gricean Model of Meaning	73
2.5 Summary	74
2.6 Exercises and Discussion Questions	75
3 Later Approaches to Implicature	77
3.1 Neo-Gricean Theory	77
3.2 Relevance Theory	91
3.3 Comparing Neo-Gricean Theory and Relevance Theory	98
3.4 Summary	107
3.5 Exercises and Discussion Questions	108
4 Reference	110
4.1 Referring Expressions	110
4.2 Deixis	114
4.3 Definiteness and Indefiniteness	121
4.4 Anaphora	130
4.5 Referential and Attributive Uses of Definite Descriptions	138
4.6 Summary	142
4.7 Exercises and Discussion Questions	143

5	Presupposition	146
5.1	Presupposition, Negation, and Entailment	146
5.2	Presupposition Triggers	152
5.3	The Projection Problem	155
5.4	Defeasibility	157
5.5	Presupposition as Common Ground	163
5.6	Accommodation	167
5.7	Summary	172
5.8	Exercises and Discussion Questions	173
6	Speech Acts	175
6.1	Performative Utterances	175
6.2	Felicity Conditions	183
6.3	Locutionary Acts	186
6.4	Direct and Indirect Speech Acts	191
6.5	Face and Politeness	200
6.6	Joint Acts	202
6.7	Summary	203
6.8	Exercises and Discussion Questions	204
7	Information Structure	207
7.1	Topic and Focus	210
7.2	Open Propositions	215
7.3	Discourse-Status and Hearer-Status	217
7.4	Information Structure and Constituent Order	219
7.5	Functional Compositionality	229
7.6	Summary	235
7.7	Exercises and Discussion Questions	236
8	Inferential Relations	241
8.1	Inferential Relations at the Constituent Level	243
8.2	Inferential Relations at the Propositional Level	260
8.3	Summary	268
8.4	Exercises and Discussion Questions	269
9	Dynamic Semantics and the Representation of Discourse	271
9.1	Theoretical Background	272
9.2	Static vs. Dynamic Approaches to Meaning	276
9.3	Discourse Representation Theory	278
9.4	The Scope of DRT and the Domain of Pragmatics	284
9.5	Summary	290
9.6	Exercises and Discussion Questions	291

10	Conclusion	293
10.1	The Semantics/Pragmatics Boundary Revisited	294
10.2	Pragmatics in the Real World	296
10.3	Pragmatics and the Future of Linguistic Theory	302
10.4	Summary	304
10.5	Exercises and Discussion Questions	304
	References	306
	Sources for Examples	314
	Index	318

Preface

Introduction to Pragmatics provides a thorough grounding in pragmatic theory for graduate students and upper-level undergraduates. While ideally the reader will come to it with a basic understanding of the principles of linguistic analysis, the text assumes little or no prior study of linguistics, and hence should be appropriate for students at all levels of expertise. In length, depth, and scope, it is suitable for a semester- or quarter-long course in linguistic pragmatics.

Pragmatics is a field that is in many ways grounded in semantics. Many of its fundamental principles have been developed in reaction to semantic principles or problems of semantic analysis; for example, Grice developed his theory of implicature in order to address the semantic analysis of the natural-language equivalents of the logical operators (such as *and* and *or*). Since its inception as a field, pragmatics has been in conversation with, and defined in opposition to, the field of semantics. The question of how pragmatics relates to, and differs from, semantics constitutes a thread running throughout this textbook. Different schools of pragmatics differ with respect to how they draw the boundary between semantics and pragmatics, a question with important ramifications for the analysis of natural language. For this reason, this question constitutes a recurring theme in this book. The text begins, therefore, with a quick review of the semantic principles and logical notation that the student will encounter in later chapters, and a discussion of the issues surrounding the demarcation of the fields of semantics and pragmatics. The text goes on to present the time-honored basic concepts of pragmatics – such as implicature, speech acts, presupposition, and deixis – while also including more recent developments in areas such as neo-Gricean pragmatics, Relevance theory, information structure, and Discourse Representation Theory.

Organization of the Book

The text consists of 10 chapters, a references section, a sources for examples section, and an index. More fundamental concepts are presented earlier, with

later chapters building on topics introduced earlier; for instance, the chapter detailing Grice's theory of implicature is followed by a chapter in which more recent approaches to implicature are discussed in light of developments over the decades since Grice's initial work on the topic. Interdisciplinary strands are woven throughout the text, as the interrelationships between pragmatics and philosophy, syntax, semantics, and even more applied fields such as law and artificial intelligence are explored. Each chapter ends with exercises and discussion questions. These are designed not only to reinforce the student's learning of the material in the chapter, but also to extend these concepts in new directions, for example by asking students to consider new variations on the chapter's theme, examine apparent counterexamples, or apply theoretical concepts to examples from their own life.

As noted above, the textbook is designed for either a quarter- or semester-long course in pragmatics at the graduate or upper-level undergraduate level. In a 9- or 10-week quarter, the instructor might choose to assign one chapter per week; in such a course, take-home exams or term papers can be assigned in order to reserve class time for discussion of the topics introduced in the text. In a semester-long course, the text can be taken at a more leisurely pace, with time available for in-class exams. For graduate courses, the text might be paired with seminal papers in each area, including primary readings from Grice, Austin, Searle, and others whose work is discussed herein; discussion of a given chapter in one class period could then be followed by a second class period in which the primary material is discussed. In this way the text would provide the necessary background for full comprehension of the primary works. Throughout, I would encourage instructors to illustrate the course material with real-life examples, both their own and those brought in by their students. Only through application to naturally occurring linguistic data can pragmatic theory be fully grasped and appreciated.

Acknowledgments

I am deeply and eternally grateful to my mentors in pragmatics – Gregory Ward, Ellen Prince, Larry Horn, and Barbara Abbott. They are my models for what a scholar should be. I am particularly grateful to Gregory Ward, who is the reason I entered the field of pragmatics and the reason I know anything at all about how to be a linguist. My debt to him is incalculable.

I wish Ellen Prince had lived to see the publication of this book, which owes so much to her outstanding work in pragmatics; her research provides the theoretical foundation that underlies all of my own. Her death has been a great loss to the field, and she is sorely missed.

For reading and commenting on early chapters of this book, I am profoundly grateful to Barbara Abbott, Larry Horn, Jeff Kaplan, Craige Roberts, Jerry Sadock, and Gregory Ward. This book has benefitted greatly from their detailed and insightful comments. I owe a special and enormous debt of gratitude to Barbara Abbott and Jeff Kaplan, who provided copious and extremely helpful comments on the entire manuscript.

I thank Andy Kehler for helpful suggestions in the early stages of this project, Larry Horn for getting me started on it in the first place, and Jeff Kaplan for making me put in the hours to finish it. I am grateful to Leah Kind for her many hours of work as a graduate assistant scouring early chapters for typos, to Nancy Hedberg for catching some embarrassing errors, and to Matt Duncan for help with the Japanese example in Chapter 7. For helpful comments and suggestions, I am grateful to many sharp-eyed students in my semantics and pragmatics classes, especially Nyssa Bulkes, Floyd Knight, Shelley Korth, Chelsea Maney, Natalie Santiago, Jessica Schlueter, and Jana Thompson. Any errors and omissions that remain are my own darned fault.

I am very, very grateful to my wonderful editor, Danielle Descoteaux, for her helpful suggestions, good humor, and preternatural patience; to Julia Kirk for keeping me organized and gently moving me along; and to Javier Kalhat for his excellent copy-editing.

I thank my colleagues and students in the English Department at Northern Illinois University for making it a pleasure to go to work every day. I am particularly grateful to five special friends and colleagues – Phil Eubanks, Kathleen Renk, Michael Day, Bonnie Anderson, and Angie Dybas – for improving my years as

Graduate Director in innumerable ways and making it possible for me to juggle the tasks of book-writing, teaching, and grad-directing.

Finally, I thank my husband Andy and my daughter Suzanne for putting up with me, serving as sounding boards, providing native-speaker intuitions, and encouraging me in this project – and for many much-appreciated work breaks playing board games at the dining room table. They have kept me sane, and that is no mean feat.

1 Defining Pragmatics

What did they mean by that? It's a relatively common question, and it's precisely the subject of the field of pragmatics. In order to know what someone meant by what they said, it's not enough to know the meanings of the words (semantics) and how they have been strung together into a sentence (syntax); we also need to know who uttered the sentence and in what context, and to be able to make inferences regarding why they said it and what they intended us to understand. *There's one piece of pizza left* can be understood as an offer ("would you like it?") or a warning ("it's mine!") or a scolding ("you didn't finish your dinner"), depending on the situation, even if the follow-up comments in parentheses are never uttered. People commonly mean quite a lot more than they say explicitly, and it's up to their addressees to figure out what additional meaning they might have intended. A psychiatrist asking a patient *Can you express deep grief?* would not be taken to be asking the patient to engage in such a display immediately, but a movie director speaking to an actor might well mean exactly that. The literal meaning is a question about an ability ("are you able to do so?"); the additional meaning is a request ("please do so") that may be inferred in some contexts but not others. The literal meaning is the domain of semantics; the "additional meaning" is the domain of pragmatics.

This chapter will largely consider the difference between these two types of meaning – the literal meaning and the intended and/or inferred meaning of an utterance. We will begin with preliminary concepts and definitions, in order to develop a shared background and vocabulary for later discussions. A section on methodology will compare the corpus-based methodology favored by much current pragmatics research with the use of introspection, informants, and experimental methods. Then, since no discussion of pragmatics can proceed without a basic understanding of semantics and the proposed theoretical bases for distinguishing between the two fields, the remainder of the chapter will be devoted to sketching the domains of semantics and pragmatics. A discussion of truth tables and truth-conditional semantics will both introduce the logical notation that will be used throughout the text and provide a jumping-off point for later discussions

of theories that challenge the truth-conditional approach to the semantics/pragmatics boundary. The discussion of the domain of semantics will be followed by a parallel discussion of the domain of pragmatics, including some of the basic tenets of pragmatic theory, such as discourse model construction and mutual beliefs. The chapter will close with a comparison of two competing models of the semantics/pragmatics boundary and an examination of some phenomena that challenge our understanding of this boundary.

1.1 Pragmatics and Natural Language

1.1.1 *Introduction and preliminary definitions*

Linguistics is the scientific study of language, and the study of linguistics typically includes, among other things, the study of our knowledge of sound systems (phonology), word structure (morphology), and sentence structure (syntax). It is also commonly pointed out that there is an important distinction to be made between our **competence** and our **performance**. Our competence is our (in principle flawless) knowledge of the rules of our own **idiolect** – our own individual internalized system of language that has a great deal in common with the idiolects of other speakers in our community but almost certainly is not identical to any of them. (For example, it's unlikely that any two speakers share the same set of lexical items.) Our performance, on the other hand, is what we actually do linguistically – including all of our hems and haws, false starts, interrupted sentences, and speech errors, as well as our frequently imperfect comprehension: Linguists commonly point to sentences like *The horse raced past the barn fell* as cases in which our competence allows us – eventually – to recognize the sentence as grammatical (having the same structure as *The men injured on the battlefield died*), even though our imperfect performance in this instance initially causes us to mis-parse the sentence. (Such sentences are known as **garden-path** sentences, since we are led “down the garden path” toward an incorrect interpretation and have to retrace our steps in order to get to the right one.)

Pragmatics may be roughly defined as the study of language use in context – as compared with semantics, which is the study of literal meaning independent of context (although these definitions will be revised below). If I'm having a hard day, I may tell you that my day has been a nightmare – but of course I don't intend you to take that literally; that is, the day hasn't in fact been something I've had a bad dream about. In this case the semantic meaning of “nightmare” (a bad dream) differs from its pragmatic meaning – that is, the meaning I intended in the context of my utterance. Given this difference, it might appear at first glance as though semantic meaning is a matter of competence, while pragmatic meaning is a matter of performance. However, our knowledge of pragmatics, like

all of our linguistic knowledge, is **rule-governed**. The bulk of this book is devoted to describing some of the principles we follow in producing and interpreting language in light of the context, our intentions, and our beliefs about our interlocutors and their intentions. Because speakers within a language community share these pragmatic principles concerning language production and interpretation in context, they constitute part of our linguistic competence, not merely matters of performance. That is to say, pragmatic knowledge is part of our knowledge of how to use language appropriately. And as with other areas of linguistic competence, our pragmatic competence is generally **implicit** – known at some level, but not usually available for explicit examination. For example, it would be difficult for most people to explain how they know that *My day was a nightmare* means that my day (like a nightmare) was very unpleasant, and not, for example, that I slept through it. Nightmares have both properties – the property of being very unpleasant and the property of being experienced by someone who is asleep – and yet only one of these properties is understood to have been intended by the speaker of the utterance *My day was a nightmare*. The study of pragmatics looks at such interpretive regularities and tries to make explicit the implicit knowledge that guides us in selecting interpretations.

Because this meaning is implicit, it can be tricky to study – and people don't even agree on what is and isn't implicit. One could make a strong argument that *a nightmare* in *My day was a nightmare* is actually quite explicit, that this metaphorical meaning has been fully incorporated into the language, and that it should be considered literal, not inferential (i.e., semantic rather than pragmatic). This in itself is a very interesting question: Every figure of speech began as a brand-new but perfectly interpretable utterance – one could say *My day was one long, painful slide down an endless sheet of coarse-grain sandpaper* – that eventually became commonplace. Upon their first utterance, such figures of speech require pragmatic inference for their interpretation; the hearer must (whether consciously or subconsciously) work out what was intended. It's possible that this is still what's done when the figure of speech becomes commonplace; it's also possible that it becomes more like a regular word, whose meaning is simply conventionally attached to that string of sounds. If the latter is the case, it's obviously impossible to say precisely when its status changed, since there was no single point at which that happened – which is to say, the shift from pragmatic meaning to semantic meaning, if and when it occurs, is a continuum rather than a point.

One might ask why it matters – but in fact there are a great many reasons why it matters. We'll return in the last chapter to some specific real-world ramifications of pragmatics, but for the present moment, just consider a court of law: It matters enormously what counts as “the truth, the whole truth, and nothing but the truth.” Does inferential meaning count as part of that truth? Courts have frequently found that for legal purposes, only literal truth matters; that is, in saying *There's one piece of pizza left*, you can be held responsible for the number of pieces of pizza left, but not for any additional meaning (such as “offer” vs.

“scolding”). On the other hand, we’ll see in Chapter 10 that the courts haven’t been entirely consistent on this issue. More generally, most people can think of cases within their own relationships in which what the speaker intended by an utterance and what the hearer took it to mean have been two entirely different things; rather sizeable arguments are sometimes due to a difference in pragmatic interpretation, with each party insisting that their interpretation constitutes what was “said.”

Pragmatics, then, has to do with a rather slippery type of meaning, one that isn’t found in dictionaries and which may vary from context to context. The same utterance will mean different things in different contexts, and will even mean different things to different people. The same noun phrase can pick out different things in the world at different times, as evidenced by the phrase *this clause* in *This clause contains five words; this clause contains four*. All of this falls under the rubric of pragmatics. In general terms, pragmatics typically has to do with meaning that is:

- non-literal,
- context-dependent,
- inferential, and/or
- not truth-conditional.

We’ll talk a lot more about that last one (“not truth-conditional”) later on; for now, it’s enough to notice that when I say *There’s one piece of pizza left*, the truth of that statement has everything to do with how many pieces of pizza are left, and nothing to do with whether I intend the statement as an offer or a scolding. Thus, the conditions under which the statement is true don’t depend on its pragmatic meaning; that’s what we mean when we say that the pragmatic meaning is generally not truth-conditional.

The “and/or” in that bulleted list is the real problem. Linguists disagree on which of these are actually defining properties of pragmatics. A prototypical case of pragmatic meaning is indeed non-literal, context-dependent, inferential, and not truth-conditional. However, there are other cases in which it’s not so clear. The case of *this clause* is a good example: Many linguists would say that determining which clause is being referred to requires a pragmatic inference, even though it affects the truth conditions of the utterance. (That is, which clause is being referred to crucially affects the question of whether *This clause contains four words* is true.) Others would say that any piece of meaning that affects truth is by definition semantic. Thus, the boundary between what counts as semantics and what counts as pragmatics is still a matter of open debate among linguists, and it will recur throughout this book as an important theme.

1.1.2 *Situating pragmatics within the discipline of linguistics*

Language use involves a relationship between **form** and **meaning**. As noted above, the study of linguistic form involves the study of a number of different

levels of linguistic units: **Phonetics** deals with individual speech sounds, **phonology** deals with how these sounds pattern systematically within a language, **morphology** deals with the structure of words, and **syntax** deals with the structure of sentences. At each level, these forms may be correlated with meaning. At the phonetic/phonological level, individual sounds are not typically meaningful in themselves. However, intonational contours are associated with certain meanings; these associations are the subject of the study of **prosody**. At the morphological level, individual words and morphemes are conventionally associated with meanings; this is the purview of **lexical semantics** and **lexical pragmatics**. And at the sentence level, certain structures are conventionally associated with certain meanings (e.g., when two true sentences are joined by *and*, as in *I like pizza and I eat it frequently*, we take the resulting conjunction to be true as well); this is the purview of **sentential semantics**. Above the level of the sentence, we are dealing with pragmatics, including meaning that is inferred based on contextual factors rather than being conventionally associated with a particular utterance.

Pragmatics is closely related to the field of **discourse analysis**. Whereas morphology restricts its purview to the individual word, and syntax focuses on individual sentences, discourse analysis studies strings of sentences produced in a connected discourse. Because pragmatics concentrates on the use of language in context, and the surrounding discourse is part of the context, the concerns of the two fields overlap significantly. Broadly speaking, however, the two differ in focus: Pragmatics uses discourse as data and seeks to draw generalizations that have predictive power concerning our linguistic competence, whereas discourse analysis focuses on the individual discourse, using the findings of pragmatic theory to shed light on how a particular set of interlocutors use and interpret language in a specific context. In short (and far too simplistically), discourse analysis may be thought of as asking the question “What’s happening in this discourse?,” whereas pragmatics asks the question “What happens in discourse?” Pragmatics draws on natural language data to develop generalizations concerning linguistic behavior, whereas discourse analysis draws on these generalizations in order to more closely investigate natural language data.

1.1.3 *Methodological considerations*

It should be noted that (like all of linguistics) the study of pragmatics is inherently **descriptive**, describing language as it is actually used, rather than **prescriptive**, prescribing how people “ought” to use it according to some standard. A linguist will never tell you not to split your infinitives; they will simply observe that people do indeed split their infinitives, and include this in their descriptive observations of language use.

Although it may seem obvious that we as scientists are interested in describing language use rather than in telling language users how they should speak, the terminology of the field can sometimes confuse the issue. For example, the Cooperative Principle to be discussed in Chapter 2 presents a series of maxims phrased

as imperatives – “say enough,” “don’t say too much,” and so on. In truth, however, these are not rules that language users are being required to follow, but rather descriptions of the principles that they typically **do** follow, and which they expect each other to follow. Nobody has to be explicitly taught to follow these guidelines; instead, they are part of what we implicitly know as speakers of our language. Therefore, it is important to keep in mind that although some of the principles described in this book are phrased in imperative form, they actually describe what speakers do automatically in using language. Rather than “speakers should do X,” what is really meant is “speakers (consistently and reliably are observed to) do X.”

In order to determine what it is that speakers do, linguists have traditionally used one of three basic methods to study language use and variation:

1. Native-speaker intuitions
 - a. Your own (introspection)
 - b. Someone else’s (informants)
 - questionnaires
 - interviews
2. Psycholinguistic experimentation
 - lexical decision, eye tracking, etc.
3. Naturally occurring data
 - a. Elicitation
 - b. Natural observation
 - c. Corpus data

The first of these, the researcher’s own **intuition**, is valuable during the initial stage of research, during the process of forming a hypothesis. It helps to guide the researcher toward a reasonable hypothesis and away from hypotheses that are clearly untenable. But once you have a hypothesis, your intuition becomes unreliable, since it may be biased toward confirming your own hypothesis. A better option is to use the intuitions of a group of informants via questionnaires or interviews, but here too you must be careful: Subjects may (consciously or not) try to please or impress you by reporting their speech as more prescriptively “correct” than it actually is. This is the “observer’s paradox” (Labov 1972): The presence of the observer affects the behavior of those being observed. Moreover, people often don’t have accurate knowledge of how they speak when they’re not paying attention.

Psycholinguistic experimentation is able to eliminate some of these difficulties by testing people’s actual linguistic knowledge and behavior outside of their ability to manipulate this behavior. For example, a lexical decision task might ask subjects to read a text and then present them with either a common word of the language or a nonsense word; their task is to determine whether the word shown is real or not. Words made salient or cognitively “accessible” by the prior text are more quickly identified as real words than are unrelated words.

Similarly, eye-tracking apparatus can determine precisely where someone is looking at a given instant (to determine, for example, what the individual takes to be the referent of a particular pronoun in a presented text, or what part of a sentence takes the longest to understand). But again, very careful set-up and control of the experiment are required in order to eliminate the observer's paradox. Typically, care is taken to ensure that the subject is unaware of what is actually being tested.

The use of **naturally occurring data** gets around these difficulties by observing language in actual use under natural conditions. **Elicitation** (in which the researcher creates a context that's conducive to getting the subject to utter the desired form) is only an improvement over intuitions if the subject is unaware that they're being observed. William Labov is famous for (among other things) a dialect study in which he asked department-store workers about the location of various items; in truth, he was merely eliciting the words *fourth floor* in order to determine which individuals dropped the [r] sound from each of the words (Labov 1966). **Natural observation** is like elicitation, except that rather than setting up a context to compel your subject to utter the desired form, you simply wait in some natural setting and watch, hoping that they will do so – and that they will do so with sufficient frequency to give you enough data to be useful. However, depending on the frequency of the desired form, one could wait quite a long time before collecting enough data to do a proper study.

The use of **corpus data** circumvents many of the above problems, in that it involves a pre-existing collection of raw language data, typically consisting of millions of words, which have been naturally produced and which can be scoured for instances of the forms under investigation. In the past, such corpora have been extremely difficult to compile, but with the computer age has come the ability to store a virtually unlimited amount of text in an easy-to-search format. The use of corpora avoids the observer's paradox, as well as sparing the researcher the trouble of waiting for a form to be produced or trying to elicit it. The use of corpus data does, however, have its own drawbacks. For example, you must take care in selecting your data sample. If your data are skewed, so will your results be. If you only look at men's speech, your results are only valid for men's speech. If you do a corpus study but use as your corpus only romance novels from the 1990s, your results will only be valid for that group of works, and you cannot generalize them to English as a whole. Less obviously, if your corpus is entirely written, it may not accurately tell you what spoken English is like. If Labov had only conducted his experiment in a single department store, he would have gotten a skewed impression of what English is like in New York City as a whole. Thus, it is important to be certain that your data are appropriate to the hypothesis that you plan to test. Second, be aware that some of the utterances encountered in corpora will contain performance errors – all those hems, haws, false starts, and so on that do not accurately reflect the language user's linguistic competence. Thus, in interpreting the results of a corpus study, researchers inevitably make reference once again to their own imperfect intuitions in order to interpret the

data they are confronted with. The best insurance is to collect as many tokens as possible, since the more data one has, the less likely it is that a performance error here or there will pose a serious threat of corrupting one's findings.

Because of the nature of the field of pragmatics, it is especially important for researchers in this field to look at spontaneous language use in a naturally occurring context. Intuitions are notoriously unreliable for pragmatic research. Some ingenious psycholinguistic studies have been devised to test pragmatic theories, but much of the current research in pragmatics is based on the study of naturally occurring data.

Finally, the type of hypothesis you are testing should be both **falsifiable** and **predictive**. To say it should be falsifiable is not the same as saying it should be false; rather, there should be some way of testing whether it is true or false, which entails that the test allow for the possibility of its being false and present a clear answer to the question, "If my claim is false, how will this test demonstrate that it's false?" For example, consider the following claims:

A discourse sometimes begins with a greeting.

A discourse typically begins with a greeting.

A discourse always begins with a greeting.

The first claim is not falsifiable, because there is no way to show that it is false (even though it's trivially easy to show that it's true). Suppose we check 100,000 discourses and find that none begins with a greeting; we will not know for sure that our claim is false, because it's always possible that the next discourse we look at will begin with a greeting and our claim will be vindicated. The second claim appears stronger, yet it too is unfalsifiable: First, the term "typically" is vague; second (and less obviously), here again we find the possibility (however unlikely) that we've just been unlucky in our selection of data and that the next 300,000 discourses will in fact begin with a greeting and will open up the possibility that our claim was correct after all. Only the third claim is falsifiable: Discovery of a single discourse that does not begin with a greeting (under some specific definition of the word "greeting") irrevocably and irrefutably falsifies our claim. Because only the third claim is falsifiable, it is also the only one of the three that constitutes an **empirical** (i.e., testable) claim. A claim is only empirical if you can imagine a circumstance that would show that it is false. And only empirical claims are scientifically interesting.

In order to be interesting, the claim must also be predictive, in the sense of being general or generalizable. That is, the claim must not simply be about a single instance of language use; instead, it must make a general claim about an entire class of uses, and therefore also predict how speakers will behave in the future. It's not interesting to present an example of a business letter and observe that it presents a problem and offers a solution, unless you can generalize this into a claim that business letters in general are constructed in such a way as to present a problem and offer a solution. Only by showing that your prag-

matic theory applies to an entire definable class of data can you argue that the knowledge that it represents constitutes part of a native speaker's linguistic competence.

1.2 The Boundary Between Semantics and Pragmatics

No discussion of pragmatics can proceed very far without a basic understanding of semantics and the proposed theoretical bases for distinguishing between the two fields. Both deal with meaning, so there is an intuitive sense in which the two fields are closely related. There is also an intuitive sense in which the two are distinct: Most people feel they have an understanding of the “literal” meaning of a word or sentence as opposed to what it might be used to convey in a certain context. Upon trying to disentangle these two types of meaning from each other, however, things get considerably more difficult. We will spend the remainder of this chapter attempting to both describe and circumscribe the domains of semantics and pragmatics, ending with a discussion of some important phenomena that challenge traditional conceptions of the boundary between the two. We will begin with a brief survey of the field of semantics and the issues with which it concerns itself.


1.2.1 *The domain of semantics*

1.2.1.1 Word meaning

Semantic meaning is typically thought of as literal meaning of the sort one would find in the dictionary. Thus, perhaps the most straightforward place to begin a discussion of semantics is in the area of word meaning. The study of word meaning is called **lexical semantics**, as opposed to **sentential semantics**, which is the study of sentence meaning (discussed below). The meaning of a word has often been described in terms of the features necessary for a thing to count as an instance of the category described by the word; for example, the meaning of the word *dog* is that set of features by which something is known to be a dog. Most word meanings are composed of more than one such feature, so that we can talk about **lexical relations** between words, by which is meant relationships of overlap (or lack thereof) in the words' semantic features. Thus, two words that overlap in all of their semantic features are said to be **synonyms**, as in the case of *car* and *automobile* or *pail* and *bucket*. **Antonyms**, on the other hand, share all of their features except for one – and on that one, they differ in choosing either opposing ends of a continuum (**gradable antonyms**, like *hot* and *cold*) or different choices from a set of exactly two options (**complementary**

antonyms, like *dead* and *alive*). Contrary to what one might expect, then, antonyms are actually very much alike: *Hot* and *cold* have a great deal in common semantically, since both are adjectives describing temperature; they differ only in which end of the temperature scale they pick out. Gradable antonyms are easy to distinguish from complementary antonyms, since gradable antonyms can be modified to represent various points on the scale: Food can be *very hot* or *somewhat hot*, and some foods can be *hotter* than others. This is not true for complementary antonyms. While it's possible to say that a party is *really dead* or that an individual is *very alive*, these are metaphorical and relatively uncommon uses; aside from very esoteric medical discussions of, perhaps, brain death vs. heart-beat, one cannot speak in any literal way of one person being more alive than another. In the case of complementary antonyms, to not be in the category described by one word is to be in the category described by the other, assuming the categories can be appropriately applied at all. That is, as long as the entity in question is the sort of thing to which terms like *alive* and *dead* may be applied (e.g., it's a rosebush or a goldfish, not a house or a coffee mug), it is necessarily either alive or dead; if it is not alive, it is necessarily dead, and vice versa. This is not the case with gradable antonyms; if one is not *cold*, it is not necessarily the case that one is *hot*. In short, gradable antonyms permit variance along a continuum, whereas complementary antonyms present an either-or situation.

Hyponymy is also a case of feature-sharing, but in this case one word (the **hyponym**) shares all of the features of another (the **superordinate**) as well as others. For example, *poodle* incorporates all of the meaning of the word *dog*, plus more. This results in a taxonomic relationship that can be drawn in tree form:



While *poodle* and *collie* are hyponyms of *dog* (their superordinate), *dog* is in turn a hyponym of *mammal*, sharing all of the semantic features of *mammal* (fur, milk production, etc.) and more. That is, a word can simultaneously be a hyponym of one word and a superordinate of another, just as *dog* is a hyponym of *mammal* while being a superordinate of *poodle*.

Homonyms result from two distinct words having the same form, as with *light* (meaning “not heavy”) and *light* (meaning “illumination”). Such a situation results in **lexical ambiguity** – that is, a case of a single lexical form having two distinct meanings. An **ambiguous** word, phrase, or sentence is simply one

that has two or more distinct meanings. Ambiguity is to be distinguished from **vagueness**, in which the boundaries of what the term applies to are indistinct. The word *pleasant* is vague, in that there's no clearly defined cut-off between what is and isn't pleasant, whereas the word *present* is ambiguous, in that it can mean, for example, either "gift" or "current time," but neither of those meanings is particularly ill-defined in its scope.

It might seem intuitively correct to describe homonyms as a single word with more than one meaning, but it's important to recognize that while only a single lexical form is involved, *light* and *light* under the different meanings described above are actually two distinct words that happen to have the same form. This situation is to be distinguished from the case of **polysemy**, in which a single word has two related meanings, as with *nickel* (the coin) and *nickel* (the metal). This is a subtle but important distinction. In the case of polysemy, the two meanings are clearly related, and the fact that the two meanings are expressed via the same lexical form is not accidental. Most dictionaries acknowledge the distinction in the way that they list words; *bat* (the mammal) and *bat* (the baseball implement) will have separate entries in recognition of their status as homonyms, while *diamond* (the geometric shape) and *diamond* (the baseball field) will be listed as subentries under a single main entry. There are, however, very tricky cases. For example, should *ruler* (a monarch) and *ruler* (a measuring stick) be considered a case of homonymy or polysemy? The answer may differ from person to person; some people recognize the relationship between the two meanings (either historically, in that measuring sticks originally used monarchs' hand and foot lengths for measurement standards, or synchronically, in that both monarchs and measuring sticks "govern" some domain), whereas others don't. If our goal in linguistics is to describe linguistic competence, that competence will vary from person to person; one person's homonymy may well be another's polysemy.

As noted above, the meaning of a word is often taken to be that set of features by which we know that the object in question is an instance of the category described by the word; thus, the meaning of the word *boy* might be composed of the features +male and -adult, and distinguished from *man*, *woman*, and *girl* by differences in these features:

	male	adult
boy	+	-
man	+	+
woman	-	+
girl	-	-

This is the approach of **componential semantics**, which attempts to boil down the meanings of words to a set of **primitive features**. But now we have a problem: What about the meaning of the word *mare*? Using only the features listed here, it will be identical to *woman*. So we'll need to add features to distinguish them – say, *equine* and *human*:

	male	adult	human	equine
boy	+	–	+	–
man	+	+	+	–
woman	–	+	+	–
girl	–	–	+	–
mare	–	+	–	+

So far, so good. But now what happens when *cow* (an adult female bovine) and *bitch* (an adult female canine) come along? Using the features listed above, they will be indistinguishable from each other; we will need to add *bovine* and *canine* as features. And no sooner will we decide that things are now in order than *sow* (adult female porcine) will come along to disturb the works, requiring yet another feature:

	male	adult	human	equine	bovine	canine	porcine
boy	+	–	+	–	–	–	–
man	+	+	+	–	–	–	–
woman	–	+	+	–	–	–	–
girl	–	–	+	–	–	–	–
mare	–	+	–	+	–	–	–
cow	–	+	–	–	+	–	–
bitch	–	+	–	–	–	+	–
sow	–	+	–	–	–	–	+

Clearly this could go on for a very long time, with a new feature required for every new species in which a female adult has a lexicalized form. Another difficulty with componential semantics is that for many lexical items, it's not at all simple to determine what the correct set of semantic features would be. For example, what are the features that constitute the meaning of the word *sandwich*? Does an object have to include two slices of bread to count as a sandwich? Apparently not, since open-face sandwiches exist. Does bread have to be involved at all? What about a pita sandwich? What about a taco? This precise question has real-world consequences: In 2006, a Massachusetts judge ruled that a burrito is not a sandwich. A Panera Bread cafe had a stipulation in its lease preventing the opening of another sandwich shop in the same shopping center. At issue was the opening of a Qdoba outlet, which sold burritos. Panera argued that a burrito is a sandwich; the judge disagreed. What set of primitive features would determine that a meat-filled pita is a sandwich while a meat-filled tortilla is not?

As an alternative to componential semantics, **fuzzy sets** offer a way of dealing with such issues. According to fuzzy set theory, the meaning of a word is a fuzzy set, that is, a set whose boundaries are indistinct, or “fuzzy.” The set contains a central member, or **prototype**, that constitutes the “best” example of the set in

question; for example, the prototypical sandwich might consist of two slices of bread with sliced meat and cheese between them, and a condiment such as mustard. Other combinations will be more or less sandwich-like depending on their resemblance to this prototype, and toward the fuzzy boundary of the set there will be cases whose membership in the class is debatable, including stuffed pitas, tacos, and burritos.

1.2.1.2 Sentence meaning

It is intuitive to think of the meaning of a sentence as the sum of its parts – that is, that determining the meaning of *Sheila won the tournament* is simply a matter of combining the meanings of the words *Sheila*, *won*, *the*, and *tournament*. And to a great extent, this is the case. A **compositional** semantics is one that takes the meaning of a sentence to be essentially the sum of its parts, in combination with a set of rules governing the way in which the meaning of the sentence is built up from the meanings of its components in light of the syntactic structures in which they are placed; that is, it's important to remember that *Mary loves frogs* does not mean the same thing as *Frogs love Mary*, and our linguistic theory must be able to explain why. Thus, the fields of syntax and semantics overlap significantly in their areas of concern.

Just as the meanings of words can overlap partially (hyponymy) or completely (synonymy) or can be in opposition (antonymy), these semantic relations have analogs at the sentence level. For instance, **redundancy** is a case of partial repetition of meaning, as in *The child plodded slowly across the yard* (where *plod* entails *slowly*) or *My female sister is very tall* (where *sister* entails *female*). As these examples illustrate, the effect of the redundancy can range from the hardly noticeable to the patently ridiculous. Notice also that hyponymy within a sentence can give rise to redundancy: *Sister* is a hyponym of *female* (i.e., *sister* includes the meaning of *female* plus more), which is what makes the sentence *my female sister is tall* redundant. Complete overlap of meaning results in **paraphrase**; for example, *My brother is older than me* is a paraphrase of *I am younger than my brother*. In this case, the paraphrase relationship is due to the lexical relationship between *older* and *younger*, but here again, the paraphrase can be due to synonymy at the lexical level: *My couch needs to be cleaned* and *My sofa needs to be cleaned* are paraphrases due to the synonymy of *couch* and *sofa*. As we will see in the next section, paraphrases are distinguished by the fact that the two sentences are true under the same set of conditions; that is, if one is true, the other is necessarily true, and if one is false, the other is necessarily false as well.

Similarly, **antonymy** at the lexical level can give rise to **anomaly** – a clash of semantic meaning – at the sentence level, as with *?The water is quite hot, and very cold*. (Throughout this text, a question mark before a sentence or clause will indicate that it is anomalous.) Not all anomaly is attributable to antonymy; consider, for example, Noam Chomsky's famous sentence *Colorless green ideas*

sleep furiously (Chomsky 1957). Here, it seems that virtually every pair of words in the sentence clash with each other: Nothing can be both green and colorless, ideas by their nature can be neither green nor colorless, ideas can neither sleep nor do anything furiously, and it is hard to imagine what it would be to sleep furiously. Thus, the sentence is wildly anomalous. Nonetheless, it is syntactically flawless, i.e. grammatical, and this was precisely Chomsky's point: He used this sentence to show that syntax and semantics are distinct, and specifically that our knowledge of the rules of syntax is autonomous – independent of the meaning of any particular sentence. The syntactic correlate of semantic anomaly is **ungrammaticality**, as in **Dog the small slept the red rug on*. (Ungrammaticality will be indicated in this text with an asterisk.)

Finally, lexical ambiguity can give rise to ambiguity at the sentence level, as with *George walked down to the bank* (where *bank* could mean “river bank” or “financial institution”). But sentences may also exhibit **structural ambiguity**, due to the existence of two distinct syntactic analyses for the sentence, as in *Jenny ate the pizza on the table*, in which either Jenny or the pizza might be on the table, depending on the structure assigned to the sentence, specifically how much of the postverbal material is taken to be part of the direct object: *Jenny ate [the pizza on the table]* vs. *Jenny ate [the pizza] on the table*.

1.2.1.3 Formal logic and truth conditions

Semantic meaning is often represented using formal notation borrowed from the study of formal logic. It's important to understand the analysis of certain English connectives in formal logic, because the seminal works in pragmatic theory take these analyses as their starting point.

First, it is useful to distinguish between **deductive** and **inductive** logic. Deductive logic involves rules for drawing necessarily valid inferences from a set of propositions. These propositions are called **premises**, and a valid inference we can draw from a set of premises is called the **conclusion**. For example:

Premises: All students love linguistics.
 Hinkelmeyer is a student.
 Conclusion: Hinkelmeyer loves linguistics.

The conclusion is **entailed** by the premises. This means that there is no situation in which the premises could be true and the conclusion false. But notice that the validity of the deduction is totally independent of the actual truth of the premises and conclusion. It could be the case, in reality, that NOT all students love linguistics, and even that Hinkelmeyer herself despises linguistics. Nonetheless, the deduction above is valid: There is no situation in which the premises could be true and the conclusion false. This is not altered by the fact that the premises themselves may not actually be true.

Inductive logic, on the other hand, is a matter of probability. Inductive inferences are not **necessarily** true, as deductive inferences are. Here's an example of an inductive inference:

Premises: The sun has risen every day of this century.
 Tomorrow will be a day of this century.
 Conclusion: The sun will rise tomorrow.

This conclusion is very likely to be true, but it is not necessarily true by virtue of the premises. That is, the fact that the sun has risen every day of this century thus far does not in itself guarantee that it will rise again tomorrow.

Formal logic concerns itself with deductive inferences – that is, with flawlessly valid inferences. It's interesting to note that scientific experiments, on the other hand, are generally designed to lead to inductive inferences – inferences that are not necessarily true. Let's say we form a hypothesis – say, that if I hold a book three feet above the floor and let go, it will fall to the floor. And let's say I perform the experiment of releasing a book from three feet above the floor 10,000 times, and each time that I let go of the book, it falls to the floor. Based on these experiments, I may confidently infer that a book held three feet above the floor and released will always drop to the floor. But notice that this is an inductive inference; it leaves open the possibility that on the 10,001st trial, the book will fail to fall to the floor. This may be unlikely, but it is a logical possibility. And indeed, if on the 10,001st trial my friend walks in and catches the falling book before it hits the floor, my hypothesis will have been falsified and will need to be revised. For this reason, the results of scientific experiments are typically reported along with a numerical value indicating the degree of confidence in the study's conclusions, expressed as a p-value: “ $p < 0.01$ ” indicates that there is a 1-in-100 chance that the conclusion is wrong, and that the results are due to chance. Put another way, this p-value indicates a 99 percent confidence in the reliability of the findings. This is one reason why it's so important that a scientific hypothesis be in principle falsifiable: Since it's impossible to confirm beyond a doubt that the claim is true (10,000 instances of dropping a book on the floor are insufficient for certainty), it is necessary to at least know what sort of circumstance would confirm that it is necessarily false (a single instance of my friend catching it as it falls).

As noted above, formal semantics employs the notation of formal logic, which it uses as a neutral, connotation-free language for expressing the meanings of **sentences**. A sentence is a sequence of words, that is, an abstract linguistic object. An **utterance** is a sentence that's produced in some actual context (whether oral, written, or signed, as in American Sign Language). There are many sentences that have never been uttered and never will be; it's quite likely, for example, that nobody has ever before uttered the sentence *My chihuahua's favorite lampshade is submerged in the lemonade*, even though it's perfectly interpretable. A **proposition** is what a sentence expresses. Thus, the sentence *I*

read the assignment today can be used to express very different propositions depending on who utters it and when. And just as a single sentence can be used to express many different propositions, a single proposition can be expressed in a variety of sentences; *Mary spoke to Jane* and *Jane was spoken to by Mary*, for example, express the same proposition.

A proposition will be true in some **possible worlds** and false in others. A possible world is precisely what it sounds like: a way that the world could have been. The idea is that the world we happen to be living in isn't the only possible world. So the proposition "all dogs are blue" happens to be false in the real world, but there's another possible world – another way the world could have happened to be – in which it's true. On the other hand, the proposition "if a dog is blue, it is blue" is true in all possible worlds. There is no possible world in which this proposition could be false; it is necessarily true. An **analytic** sentence is one whose truth is independent of what the world is like; it's either necessarily true (as in *if a dog is blue, it is blue*) or necessarily false (as in *if a dog is blue, it is not blue*). A sentence that is true in all possible worlds (such as *if a dog is blue, it is blue*) is a **tautology**. A sentence that is not true in any possible world (such as *if a dog is blue, it is not blue*) is a **contradiction**. A sentence whose truth depends on the condition of the world (such as *some dogs are blue*) is **synthetic**. In order to know whether a synthetic sentence is true in a given world, it is necessary to see what that world is like (for example, whether it contains any blue dogs).

The **truth conditions** of a sentence are the conditions under which it would be true – that is, what the world would have to be like in order for that sentence to be true. The truth conditions of a sentence are independent of what the world actually **is** like; they're just a specification of what the world **would** be like if the sentence were true. On the other hand, the **truth value** of a sentence in some particular world is a specification of whether the sentence is in fact true in that world. Thus, the truth conditions of the sentence *A blue dog exists* are essentially that the world contains a blue dog, while the truth value of the sentence is T (true) in a world that does contain a blue dog and F (false) in a world that does not. **Truth-conditional** meaning is any piece of meaning that affects the conditions under which a sentence would be true. Thus, the difference between *and* and *or* is truth-conditional, since the sentences in (1) and (2) are true in different sets of circumstances:

- (1) All women are tall and all women are smart.
- (2) All women are tall or all women are smart.

In a world in which all women are smart but not all women are tall, (1) would be false while (2) would be true. However, the difference between *moreover* and *nonetheless* is not truth-conditional:

- (3) All women are tall; moreover, all women are smart.
- (4) All women are tall; nonetheless, all women are smart.

The sentences in (3) and (4) will be true under the same set of circumstances; there is no possible world in which one of them is true and the other false. There is, of course, an additional piece of meaning that's conveyed in (4); here you understand the speaker to be suggesting that in the context of all women being tall, there is something unexpected about their also being smart. By saying that this piece of meaning is non-truth-conditional, we don't mean that the sentence *There is something unexpected about all women being smart* has no truth conditions; it obviously does. Rather, we mean that its truth conditions play no role in determining the truth conditions of (4), and likewise that its truth value (i.e., whether it is in fact the case that this is unexpected) plays no role in determining the truth value of (4) when it's uttered.

The study of logical relationships between sentences is called **propositional calculus**. In propositional calculus, p , q , and r stand for propositions, and they are connected by various **logical connectives** such as *and* and *or*. The logical connectives can be viewed as functions that map truth values (or sets of truth values) onto truth values. For example, take logical negation:

p	$\sim p$
t	f
f	t

This is called a **truth table**. What it tells us is that anytime p is true, $\sim p$ ("not- p ") is false, and anytime p is false, $\sim p$ is true. Thus, negation is a function that maps **t** in the first column onto **f** in the second, and vice versa. In each row, the values to the left of the double line give us the truth value(s) of the given proposition(s) in some world, and the values to the right of the double line tell us what that means for the values of the propositions in combination with the given connectives. In the little truth table above, for example, the first line represents any world in which p is true; in such a world, $\sim p$ is necessarily false. The second line represents any world in which p is false; in such a world, $\sim p$ is necessarily true. Thus, if *All fish have fins* is true, then *Not all fish have fins* must be false, and vice versa. While negation isn't technically a connective (since it doesn't connect two propositions), it is typically grouped with the logical connectives because, like the logical connectives, its meaning is defined as a function from truth values to truth values. Notice that it doesn't matter what the proposition in question (p) is; the effect of negation will be the same regardless of the particular meaning of p .

The truth table for conjunction ("and," symbolized $\&$ or \wedge) is slightly more complicated, since it involves two propositions:

p	q	$p \& q$
t	t	t
t	f	f
f	t	f
f	f	f

What this table tells us is that $p \& q$ is only true when both p and q are true (the first line). In all other cases, $p \& q$ is false. That is to say, *All monkeys are mean and all fish have fins* is false if either *All monkeys are mean* is false or *all fish have fins* is false, regardless of the truth of the other conjunct.

Here's the truth table for disjunction ("or," symbolized \vee):

p	q	$p \vee q$
t	t	t
t	f	t
f	t	t
f	f	f

What this table tells us is that $p \vee q$ is false only when both p and q are false (the fourth line); in all other cases, it's true. This is the truth table for what's known as **inclusive "or,"** meaning "one or the other or both." On this reading of "or," *All monkeys are mean or all buffalo are brave* is true if either all monkeys are mean or all buffalo are brave, regardless of the truth of the other conjunct.

The truth table for **exclusive "or,"** meaning "one or the other, but not both," would be:

p	q	$p \vee q$
t	t	f
t	f	t
f	t	t
f	f	f

Here, if both propositions are true, the entire disjunction is false (line 1). This would be the meaning generally intended in the utterance of a sentence such as *I'll pay you tomorrow or the day after* (where the speaker doesn't intend to leave open the possibility of paying on both days). Exclusive "or" is usually assumed to be derived via a pragmatic inference; that is, truth-conditionally "or" is assumed to have only the inclusive meaning, but in many contexts hearers infer that it's not the case that both conjuncts are true, because if they were (and if the speaker knew they were), the speaker should have used "and."

Here's the truth table for logical implication (aka the conditional, or "if . . . then," symbolized \rightarrow):

p	q	$p \rightarrow q$
t	t	t
t	f	f
f	t	t
f	f	t