

The background of the cover is a deep black, overlaid with vibrant, glowing blue lines that swirl and curve in a dynamic, almost organic fashion. These lines vary in thickness and intensity, creating a sense of movement and depth. The overall effect is reminiscent of a stylized representation of neural pathways or the complex structure of language and cognition.

Cedric Boeckx

Language in Cognition

Uncovering Mental Structures
and the Rules Behind Them

 WILEY-BLACKWELL

Praise for *Language in Cognition*

“Boeckx has a deep familiarity with all of the (very wide-ranging) material he presents, and has done original and important work in several of these areas. He is a lucid and engaging expositor, and is highly qualified in every respect to undertake an enterprise of this nature . . . [He] brings together the right topics, some right at the edge or even at the horizons of research. If I were teaching undergraduate or graduate courses in these areas, I cannot think of a competing text that I would prefer.”

*Noam Chomsky, Massachusetts Institute of Technology,
Professor Emeritus of Linguistics*

“Boeckx masterfully presents the material, showing why biology must form a core part of the mind sciences, and how the mind sciences, and especially language, can pose new challenges for biology. It is an argument that every serious student of the mind sciences should know, even if they don’t join the choir.”

Marc Hauser, Harvard University

“No one had integrated, yet, the key four notions ‘language’, ‘cognition’, ‘mental’, and ‘structures’ in such a masterly and original way. A wealth of discoveries awaits both the novice reader and the expert. Cedric’s art of revealing deep connections between fields, authors, and ideas has frequently prompted me, while reading this book, to wonder why I had not seen those connections before.”

Massimo Piattelli-Palmarini, University of Arizona

“If Linguistics is the queen of the cognitive sciences, then Cedric Boeckx is her official court biographer. In one short book, he clearly outlines how the developments in linguistics have upended earlier empiricist conceptions of mind and spurred exciting investigations of human evolution.”

Norbert Hornstein, University of Maryland

“Cedric Boeckx provides a wonderful, modern review of the necessity of mentalism, of innate structure for all of the mind, and the role of mathematics in articulating different principles of representation for different modules of mind: a summary of the Chomskyan revolution over the last half century. He brings perspective to the project by connecting the history of philosophy with modern experimentation showing that the ‘generative’ approach to both language and mind has received stunning support in acquisition, processing, and aphasia. It is a superb introduction to the fundamental role of generative thought in modern cognitive science, weaving together psychological, biological, and philosophical perspectives, while acknowledging that fundamental aspects of human nature remain mysterious.”

Tom Roeper, University of Massachusetts Amherst

Language in Cognition

**Uncovering Mental Structures
and the Rules Behind Them**

Cedric Boeckx



WILEY-BLACKWELL

A John Wiley & Sons, Ltd., Publication

This edition first published 2010
© 2010 Cedric Boeckx

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,
United Kingdom

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 Cedric Boeckx to be identified as the author of this work has been asserted in accordance with the 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

Boeckx, Cedric.

Language in cognition : uncovering mental structures and the rules behind them / Cedric Boeckx.
p. cm.

Includes bibliographical references and index.

ISBN 978-1-4051-5881-7 (hardcover : alk. paper) — ISBN 978-1-4051-5882-4 (pbk. : alk. paper)

1. Psycholinguistics. 2. Cognitive grammar. I. Title.

P37.B64 2009
410.1'9—dc22

2009023108

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

Set in 10.5/13pt Minion by Graphicraft Limited, Hong Kong
Printed in Singapore

For Youngmi,

The light at the end of all my tunnels

Contents

Acknowledgments	viii
Prologue	xi
Part I Ever Since Chomsky	1
1 Mind Matters: Chomsky's Dangerous Idea	3
2 The Mechanization of the Mind Picture	14
3 How the Mind Grows: From Meno to Noam	41
Part II Unweaving the Sentence	55
4 Mental Chemistry	57
5 The Variety of Linguistic Experience: The Towers of Babel and Pisa	82
6 All Roads Lead to Universal Grammar	95
Part III The Mental Foundations of Behavior	109
7 Making Sense of Meaning: An Instruction Manual	111
8 Wonderful Mental Life: Unthinkable without Language	120
9 Grammar Caught in the Act	133
Part IV Missing Links	147
10 The (Mis)Measure of Mind	149
11 <i>Homo Combinans</i>	163
12 Computational Organology	173
Epilogue	183
Notes	189
Guide to Further Study	211
References	219
Index	242

Acknowledgments

The seeds of this book were sown when I joined the department of Linguistics at Harvard University in the Summer of 2003. At the time (much to my surprise) linguistics was not part of the massive Mind-Brain-Behavior initiative on campus and lacked a course fully dedicated to how the study of language bears on cognition. Thankfully, I was given the chance to introduce a new course “Language and Cognition” and was allowed to form ties with psychologists and other members of the Mind-Brain-Behavior community, which eventually resulted in the creation of a (long-overdue) Linguistics/Mind-Brain-Behavior study track. Both sets of interactions, with faculty and students, led me to think deeply about the role linguists could play in the bigger cognitive science community and ultimately helped me structure the material of the present book. First thanks therefore must go to all my students, especially those that took my pilot class on language and cognition in the fall of 2003, and to the members of the Mind-Brain-Behavior community for the warm welcome they offered me, particularly John Dowling, David Haig, Liz Spelke, and Marc Hauser. I am also grateful to Martin Nowak for making me an associate member of his center at the interface of mathematics and biology, through which I learned more about how biologists approach certain problems.

I was fortunate to be asked by Ada Brunstein to turn my thoughts on language and cognition into a book. Although she left the Blackwell team before I could finish a draft of the volume, her support from the very beginning was very inspirational, and I am grateful to her for keeping track of my progress even after she had moved on, and, more generally, for her confidence in the project. The team at Wiley-Blackwell (Danielle Descoteaux and Julia Kirk) worked with me closely to make this a better book, and I thank them, and the authors of the anonymous reviews they commissioned, for all their efforts on behalf of this project. Thanks also to Anna Oxbury for excellent copy-editing.

For several years now I have had the good fortune to be a frequent visitor and close collaborator of various members of the department of linguistics at the University

of Maryland. Since its conception (thanks to David Lightfoot) the department has retained its unique focus on the language faculty and its mentalist agenda, offering the sort of oasis that I needed at various points. Readers of this book will quickly realize how much I have been impressed by the works of my College Park friends and their students. Special thanks go to Juan Uriagereka, Norbert Hornstein, Paul Pietroski, Howard Lasnik, Colin Phillips, David Poeppel, Jeff Lidz, Bill Idsardi, David Lightfoot, Stephen Crain, and Rozz Thornton for being so inspirational and supportive.

Thanks must go to Randy Gallistel, Alec Marantz, Gary Marcus, Steven Pinker, and Ray Jackendoff for discussing various issues of cognitive science with me at various points. Ray's work, in particular, has been a source of motivation, as it always seeks to bridge disciplinary boundaries. Although Ray and I disagree on many issues surrounding the architecture of the language faculty, his 1994 book *Patterns in the Mind* was the book I originally relied on in my language and cognition class.

My teacher in graduate school, Howard Lasnik, was the first to make me realize how important and relevant chapter 1 of Chomsky's *Aspects of the Theory of Syntax* remains 40 years after it was written, and I am extremely grateful for this, as I am for the teaching example he set for me and many others.

Noam Chomsky has been my mentor since day one. Words cannot adequately characterize what this has meant to me. All I can say is that his works were the reason I joined this field, and his constant encouragement has been the main reason I believe I can contribute to it. As for all my projects, Noam kindly listened when I sketched what I wanted to put in this book, and offered valuable advice at every step. I hope the reader will realize how much poorer linguistics and cognitive science as a whole would be without his insights. Morris Halle also deserves my thanks for his unfailing support over the years.

A host of graduate students have provided much needed help. I am grateful to Bridget Samuels, Dennis Ott, Hiroki Narita, Beste Kamali, Suleyman Ulutas, and Terje Lohndal for comments on the penultimate draft. I am especially grateful to Bridget for her editorial help at various stages and to Hiroki for compiling the index.

Special thanks go to So-One Hwang, who was one of the students in my pilot class on language and cognition, and went on to study language and cognitive science in graduate school despite pressure from various sides to settle on medical school or other "safer" career options. Her brave decision was a key motivation for me to write this book, in the hope that making this material available will convince others to do just what she did.

Harvard College provided me with various funds to help me gather material and try out various ways of conveying it to students, which helped tremendously in the completion of this project. Thanks go especially to Georgene Herschbach, Associate Dean of Harvard College for Undergraduate Academic Programs, for her support.

Over the years my friends Adam Szczegielniak, Norbert Hornstein, and Massimo Piattelli-Palmarini have spent countless hours talking with me about linguistics and cognitive science, teaching me many things I did not know, and giving me better ways to convey some of the central aspects of our field. They had a major influence

(all too insufficiently and inadequately acknowledged in the text) on this book, for which I am indebted to them.

My final set of thanks goes to the most important force behind this project, and indeed behind everything I undertake in life. My wife Youngmi saw well before me that I could write this book, and made sure I brought the volume to completion. She has provided a perfect mixture of love and support to make the writing of this book as enjoyable as any author could hope for. For this and the countless other things she does for me, I dedicate this book to her, in loving gratitude.

Prologue

Serious fields of inquiry resist simple characterizations and ready-made definitions. Cognitive science as a whole, including all its subfields, such as linguistics, is no exception. Academic terms can also be quite misleading: one should not assume that linguistics deals with language just because the two terms are etymologically related; after all, who still thinks that geometry deals with land-measurement (the literal translation of the Greek term *geometria*)?

At its most general level of description, the enterprise we call cognitive science is a massive effort to construct a scientific understanding of mental life,¹ the product of the brain – arguably the most complex object in the known universe.² Although some of the leading ideas reach back to the rise of modern science, the time of Galileo, Descartes and Newton (some ideas even go further back to the Ancient Greeks), the efforts to construct a genuine scientific theory of mental life began in earnest only 50 years ago under the impetus of people like Noam Chomsky, Morris Halle, George Miller, and Eric Lenneberg. Not surprisingly, after only a few decades of intensive research, our ignorance overall remains quite profound, but there are a few areas where significant progress has already been made. One such area concerns our capacity to develop a language, and this is the area I will focus on here, touching on other cognitive domains whenever the opportunity arises. The results achieved in the domain of language have been made possible by the adoption of what can be called the biological view of language, in which the problem of making sense of our human capacity to acquire and use a language is conceived of as being on a par with how scientists would study echo-location in bats, the waggle dance in bees, and the navigational skills of birds.

The biological approach to language will enable me to paint one of the most interesting, insightful, and coherent pictures in cognitive science we currently have, and place it right at the heart of one of the best-articulated theories of mind (and because language is unique to our species, of what it means to be human) ever produced.

I should state right away, to avoid bitter disappointments in my readers, that we will treat language (and other cognitive capacities of ours, like music, mathematics, vision, etc.) as a natural object, fit for scientific inquiry. We will not focus on all the ways in which we use language (and the rest of our minds) in daily interactions with others. Instead we will focus on our unconscious knowledge of language, that which gives us the ability to form an infinite range of expressions, and at the same time exclude countless other ways of expressing ourselves. This focus of investigation turns out to be the only way to make progress. If you know a little bit about the history of the more established sciences (physics, chemistry, biology), this should not surprise you. The way science progresses is by first acknowledging the complexity of the world, and immediately admitting the futility of attempts to provide a full description of it. Once humbled by this recognition of the vastness of the problem, the best way forward for the scientist is to extract a small corner of the problem, make a few simplifying assumptions, and attempt to gain some understanding of that small part of the world. Science – like everything else, I guess – works best when it deals with simple things. It's no use to ask questions we will never answer, even if these are the first questions that come to our ever-curious minds. Becoming a scientist is a subtle development that requires shedding our childish stubbornness to only ask questions we want to know the answers to, while at the same time preserving the childlike attitude that lets us ask questions that most adults find uninteresting.

My main goal in this book is to awaken your curiosity by pointing out a few facts that I suspect you never thought about, by asking questions that will cultivate your sense of wonder, and by suggesting a few answers that will whet your intellectual appetite. For this reason I will put less emphasis on results achieved in cognitive science and focus on the questions that have proven fruitful in making these results within reach. This book should definitely not be seen as providing a sum of all we know about the mind; it is best characterized as offering a point of entry into fascinating territory, a set of perspectives from which to approach certain topics. If you are like me, you will find some of the questions cognitive scientists ask simply irresistible. By the end of the book, you will be able to turn to more advanced material that will explore these questions in greater detail. This, I find, is the most appropriate way to introduce students to a scientific discipline. Science, by its nature, involves a constantly changing and developing body of knowledge. Today, perhaps more than ever, that knowledge develops and changes very rapidly. Because of this fact, many recent educational initiatives have stressed the need for science teachers to develop ways to help students understand the practice of scientific inquiry, and not just its current results, which may well be outdated by the time an introductory text hits the bookstores' shelves, especially in the case of young scientific disciplines like cognitive science. As educators, we want students to avoid falling into the trap of a passive dependence on "experts," and we want them to develop a critical mind. To do this, it is imperative that they come to understand how scientific knowledge is acquired, and how to derive it themselves.

The skills that everyone deems important in the conduct of scientific inquiry (pattern recognition, what-if reasoning, hypothesis-testing, and so on) are largely independent of content, in the sense that they can be illustrated using examples from any scientific domain, but I agree with many that some domains appear to offer advantages as educational media. As a few cognitive scientists have pointed out,³ the study of human language as a biological object is in many respects a great instrument for the teaching of science.

Linguist Richard Larson⁴ has summarized the main arguments for why this is the case, and I want to reproduce them in his terms here:

- 1 There is the sheer accessibility of the subject matter, especially data; language is all around us, so data can be collected easily. At least the basic data does not require high-level mathematics or complex tools. As a result, students can move very quickly from data collection to scientific questions and analyses.
- 2 There is the intrinsic interest of the subject matter. Every popular book about linguistics starts off by pointing out that everyone is interested in language, and rightly so. As we will see later on in this book, there is good evidence that children are biologically built in such a way that they can't avoid paying attention to language; our biology makes us, in some sense, natural linguists. People love to manipulate language, look at different ways of expressing various thoughts, coin new terms, etc. Once carefully monitored by the teacher, this natural interest can turn into real scientific investigation at no cost.
- 3 There are certain social factors that contribute to make linguistics an ideal point of entry into science. As Larson notes, the fact that its data can be easily gathered largely means that linguistics is fully accessible to students with physical limitations that might otherwise present a significant barrier to learning (blindness, gross motor impairment, etc.). Furthermore, I think Larson is right that linguistics is a "gender neutral" domain. Being a young discipline, linguistics lacks some of the stereotypes that define who is a canonical "language scientist."
- 4 Linguistics is a natural "bridge" discipline between the sciences and the humanities. Because its subject matter – language – is at the very center of many areas in the humanities, linguistics offers an excellent way to appeal to students not otherwise considering science as an area of study, and to introduce them to the principles of scientific method and reasoning.

In addition, as linguist Ray Jackendoff has observed,⁵ the reasons why many schools virtually gave up on teaching grammar no longer apply when language is studied as part of cognition. As we will see in subsequent chapters, modern linguistics as a sub-branch of cognition moves away from prescriptive grammar and its arbitrary rules of "proper speech"; it provides much better descriptions than traditional non-starters like "nouns are names for persons, places, or things" (how about *nothing* or *redness*?), it recognizes socially disadvantaged students' own ways of speaking as valid subjects for grammatical description (African American English is just as

good a medium to probe cognition as any other English variety), it highlights the importance of studying other languages, it makes short shrift of the claim that children will get confused if they try to speak more than one language (look around the world: multilingualism is the norm rather than the exception), and it provides a good illustration of the well-foundedness of and the subtle questions raised by evolutionary theory in biology.

For all these reasons (illustrated in detail throughout this book), I think cognitive science, with linguistics at its center, ought to be recognized as an essential part of the curriculum.

Although the book is primarily aimed at students with little or no background in the subject matter, I have provided extensive references in the Notes section at the end of this book, which I hope will make it possible for the book to function as a useful guide for more advanced students who may want a “big picture” presentation, so easily obscured by technical aspects of the various subfields involved (aspects that are, of course, necessary, but never sufficient). I certainly encourage curious students to track down as many of the references I provide as they can. I also encourage them, as well as the instructors who adopt this book for courses, to use the study guide at the end of this book and assign readings listed there as supplements, like fine wine is supposed to accompany a good meal. The users of this book should find enough flexibility in the material presented here to suit various needs, and cater to a variety of audiences.

To conclude this brief overview let me emphasize that in presenting the material that follows I make absolutely no claim to originality. The exciting discoveries reported on here, the guiding intuitions, the questions asked, etc., are not mine. All the credit should go to the authors mentioned in the bibliography. My task was merely organizational and journalistic: I have selected material from various disciplines that seem to me to shed light on our linguistic capacity and the nature of cognition, and to offer complementary methods to investigate a coherent set of questions. It stands to reason that in selecting material I am offering a rather personal view of what cognitive science is and what cognition may be. I am presenting what I take to be our most promising hypothesis regarding mental life, but I have no doubt that others would have made different choices. Throughout I have favored intellectual coherence over exhaustive coverage, and often relegated alternatives to endnotes. I encourage readers to track down these competing viewpoints, make sense of them, and improve on the perspective presented here.

Because my task was merely one of selection, collection, and organization, I have incurred many intellectual debts, and I am grateful to the authors of all the works cited in the bibliography for giving me such an abundance of excellence. I am especially indebted to Noam Chomsky for providing the guiding ideas and some of the best results in this book. Yes, this book is unabashedly Chomskyan in character, in the same way that introductions to evolutionary biology must be unabashedly Darwinian. I have been impressed by how much can be gained by asking the sort of questions Noam Chomsky asks, and I hope the reader will be equally impressed and motivated to ask them. I find the Chomskyan line of

inquiry the best antidote against our denial of human nature and against our own arrogance. I am also indebted to Alec Marantz, Colin Phillips, David Poeppel, Marc Hauser, Tecumseh Fitch, Liz Spelke, Paul Pietroski, Juan Uriagereka, and Massimo Piattelli-Palmarini, whose works have provided much of the ideas presented in Parts III and IV of this book, the most interdisciplinary parts of the volume. The works of the late David Marr, Jerry Fodor, Randy Gallistel, Charles Yang, Ray Jackendoff, and Gary Marcus have also been extremely valuable in preparing various chapters of this book.

I should also point out that although the book ranges over a broad spectrum of ideas, I am ultimately a tradesman, not a polymath. What I know about neuroscience, biology, and cognition in general is what lies at their intersection with human language. If some readers of the book feel that I have failed to report on some discoveries that they deem essential, I apologize in advance.

Throughout the book, I focus mainly on issues involving syntax, the study of sentence structure. This should not be taken to imply that syntax is more important or more central than other areas of language like phonology (sound structure) or morphology (word structure). It merely reflects the fact that I am first and foremost a syntactician, and also the fact that syntax is one area of language where enough of the pieces are in place to allow us to seriously consider how they all fit together, and to identify what is still missing. No doubt, other subfields of linguistics would have served equally well for purposes of illustration at various points in the book, and instructors that are less enamored with syntax than I am should feel free to substitute their own favorite topics.

With all these warnings in place, I now ask the reader to turn the page and reflect on the problems and mysteries of our own human condition.

PART I

Ever Since Chomsky

1

Mind Matters: Chomsky's Dangerous Idea

1 Introduction

If you have ever crossed Harvard Yard, as I have done repeatedly on my way to teach the class on which this book is based, you must have seen visitors gathered around the statue of John Harvard. The statue is known as “The statue of the three lies.”¹ The inscription on the stone supporting the statue reads: “John Harvard, Founder, 1638.” In fact, the model for the statue used by the sculptor (Daniel Chester French) was a student, since there was no portrait of John Harvard available. Equally misleading is the title “founder.” Harvard College was not founded by John Harvard; it was named after him. And the College was founded in 1636, not 1638.

It must have been on one of those occasions when I passed by the statue that I toyed with the idea of calling this book “Chomsky’s Dangerous Idea.” As is the case with the Harvard statue (or, for that matter, with terms like the Holy Roman Empire, which, as Voltaire once shrewdly remarked,² was neither holy, nor Roman, nor an empire), the phrase “Chomsky’s dangerous idea” wouldn’t be accurate, but would nonetheless capture the essence of what I am trying to do in this volume. Although the book will not deal with just one idea, and it won’t always be Chomsky’s, and – I assure you – none of the ideas discussed here are dangerous, it is nevertheless quite correct to say that the central argument developed in the following chapters, shorn of all the qualifications and additions that many researchers have brought forward, goes back to Noam Chomsky. Furthermore, Chomsky’s central claim regarding language remains, in the eyes of many, quite controversial.

What’s this central idea? It will take me a book to characterize it properly, but for now, let me try to state it in simple terms. We, as human beings, come equipped biologically with the ability to develop a language, to make linguistic sense of the noise around us. This ability is both extremely rich in its potential and severely constrained, in ways that cannot be understood by simply looking at properties of

the world around us in which we use language. It requires a *mentalist* stance: a readiness to posit principles of the mind that we, as humans, bring to the task of language acquisition and language use. Since the very beginning of his career,³ Chomsky has urged us to look at the mind in a way parallel to how biologists study the heart, the kidney, or the immune system: as a complex organ of the body, where nature and nurture interact to give rise to some of our most distinctive traits as a species. (For this reason, the approach to language Chomsky advocates is often, and quite appropriately, called *biolinguistics*.⁴)

“Chomsky’s dangerous idea” could also serve as an allusion to Dan Dennett’s popular book entitled *Darwin’s dangerous idea*.⁵ This, too, would be apt, as I happen to think that there are quite a few interesting parallelisms between Chomsky and Darwin. Few scientists have had such a lasting influence on their respective disciplines as these two. Both are extremely controversial figures: They are either reviled or revered (much to their dismay). Both reoriented their own respective fields by articulating a vision that promises to yield genuine scientific understanding and shed light on some of the most difficult questions ever raised. Both showed an impressive ability to revise old ideas, and to turn these ideas into testable scientific hypotheses. Both have advocated theories that, upon reflection, make eminent sense. Indeed, many a scientist would regard Chomsky’s view on language and Darwin’s view on life as virtually the only game in town – the very best we have in these domains of inquiry. And yet, in spite of, or perhaps because of the simplicity of their views, quite a few researchers resist them, and they still require lengthy defenses and expositions.⁶ Finally, like Darwin, the implications of Chomsky’s views for human nature, though implicit in the work that made him famous, only became clear in subsequent publications.⁷

In the end I decided against using “Chomsky’s Dangerous Idea” as the title for this book. I wanted something less personal, and more immediately transparent, something with key terms like “language,” “cognition,” “mental,” and “structures,” to serve as *leitmotifs*. These will be the real characters in my story. My overall aim is to convince you that the study of human language offers a unique vantage point from which to understand human nature (“languages are the best mirrors to the mind,” as Leibniz once asserted⁸), from which to do cognitive science, from which to dig deep into what William James called “the science of mental life.”⁹ Linguistics, as practiced along Chomskyan lines, is an essential player in the elaboration of a genuine “Science of Man.”¹⁰

I would be happy if by the end of the book I have managed to convey to the reader some feeling for what is without a doubt one of the most important and exciting paths of discovery that our species has embarked upon: the search for the underlying principles that both govern and constrain what our brain does (i.e., what our mind is). It is an avenue of research that many have explored, but where progress has been very slow. Cognitive science as we know it today is a relatively young science. It is only in the past 50 years that substantial advances have been made, and much – so much! – remains to be discovered. To the novice, this can be both a source of frustration, and excitement. Frustration because even some

of the most basic questions are still up for grabs, leading to the impression that the foundations are not firm, that too much is changing too rapidly, that in the end very little is known, and that much has to be unlearned on a frequent basis. (These are impressions that I have invariably encountered in my teaching, and that I will try to dispel. I'm introducing them now to urge the reader to recognize them and resist them.)

Excitement because even some of the most basic questions are still up for grabs, leading to the impression that fundamental discovery could come from one's own research, that the missing link is within reach, and that this is history very much in the making. (These are impressions that I have invariably encountered in my teaching too, and that I will do my best to cultivate throughout our journey. May the reader never abandon them in the face of difficulty!)

2 I-language

Using language to study the human mind is a very old practice. Since at least Descartes¹¹ in the modern era (but no doubt, the practice goes back much further), it has been standard to treat language as providing privileged access to some of the deepest secrets of the mind. Unfortunately, the view of language that is required to probe cognition is quite different from our common-sense notion of language, and the way we experience it on a daily basis. We tend to think of language as something that is hard to acquire, that varies from place to place, that is inextricably linked to social norms and culture, and so on. And yet I will argue, following Chomsky, that if we are to make progress in the domain of cognition, language must be understood as something that is acquired effortlessly, that is shared by the entire human species and fundamentally the same across cultures, and that is radically dissociated from social norms. This, I realize, will be hard to swallow at first, and part of this chapter will be devoted to rendering this view somewhat plausible. In an attempt to make this change of perspective conspicuous, Chomsky has suggested¹² we distinguish between language as seen from a social/cultural perspective (what he calls "*E(xternal)-language*"), and language as seen from a cognitive perspective (what he calls "*I(nternal)-language*"). Unless otherwise indicated, whenever I use language, I mean "I-language." (This is what the word "in" in the title of the book, *Language in Cognition*, is meant to emphasize.)

I ask the reader to keep the distinction between I-language and E-language firmly in mind because invariably, when I find myself in a casual setting and I mention the fact that I study linguistics, I am asked how many languages I speak. To a linguist like myself, this question is distinctly odd, for many of us (myself included) think there is only one human language on the planet. This is not to deny that there exist many linguistic variants (languages, dialects, etc.), but these are not the primary objects of inquiry to the linguist/cognitive scientist. Asking a linguist how many languages she speaks/knows is a bit like asking how many

numbers a mathematician knows, or how many species a biologist knows, or how many molecules a chemist knows.

A linguist, to the extent that she is interested in finding out how some aspects of the mind work, tries to elucidate the mental capacities in virtue of which humans are able to produce and understand utterances in *any* (human) language variant. Linguistics, in this sense, is like trying to discover the laws of (some aspects of) the mind. Like any other scientist, the linguist will have to learn to go beyond the multitude of languages to get at the fundamental principles that make linguistic experience possible. From this perspective, specific human languages (English, Navajo, Japanese, . . .) are used as tools, convenient entry points to study something more fundamental, the same way biologists study rats, fruitflies, and worms. Here the linguist has the distinct advantage that languages are literally all over. We are surrounded by language(s). For this reason, linguistic facts often appear to be “cheap.” There is no need to deal with dangerous equipment, flammable substances, etc., to get something to work on. But linguistics is not blind data collection. It is an empirical science. One must be as careful with linguistic data collection and interpretation as in a science lab. And herein lies the rub. Many of the data that linguists deal with pertain to aspects of language that we use instinctively, and that we therefore are unaware of. Grammar classes in high school don’t dwell on them, and they are not discussed in popular newspaper columns about language. And because our mastery of these facts is so effortless, we tend to think that the explanation for these facts must be straightforward and intuitive. But as George Gershwin put it, it ain’t necessarily so.

3 A Few Illustrations

Consider the fact that you can readily recognize ambiguous sentences like these (often, contextual information, or special intonation favors one reading, but given enough time all native speakers are able to detect the ambiguity):

- (1) Flying planes can be dangerous.¹³ (either planes that fly are dangerous, or flying these planes is dangerous)

or

- (2) Mary hit the man with the umbrella.¹⁴ (either the man had the umbrella, or Mary used the umbrella to hit the man)

or

- (3) John can ride a bike. (either John is allowed to, or he has the ability to, ride a bike)

Data like these are so mundane that the richness they contain hardly gets noticed. For instance, did you ever pause to marvel at the fact that a sentence like (4) is only two-way ambiguous, and not four-way ambiguous?

(4) John can ride a bike, and Mary can ride a bike, too.

On the face of it, you might have expected that putting two two-way ambiguous sentences together (*John can ride a bike; Mary can ride a bike*) would give rise to a four-way ambiguity, but somehow, it doesn't. Something forces you to understand *Mary can ride a bike* in the same way you understand *John can ride a bike* when the two sentences are combined into one. And no context, no matter how contrived, will make this fact go away. You must interpret members of a conjunction in parallel. Let's call this the Parallelism Constraint.

If you are like my students, you may be tempted to say that the Parallelism Constraint is due to the presence of the words *and* and *too*. But how about (5)?

(5) John can ride a bike, but Mary cannot ride a bike.

Here too, you have to understand the meaning of the word *can* in the same way in both sentences. But it's not clear what word to blame this constraint on this time. I'm sure you have never thought about this fact before (unless you met a linguist at a cocktail party, and he or she tried to impress!), and you were never told how to interpret ambiguous sentences. Facts like those just discussed constitute *tacit* (unconscious) knowledge: knowledge that the great linguist Morris Halle aptly called "unlearned and untaught."¹⁵ It's just part of what it means to be a native speaker of a language. (Here and so often elsewhere in the book, I will use English examples to illustrate the points I want to make, but this is just a matter of convenience. Similar examples could be found in any other language of the world.)

The extent of your tacit knowledge of language does not stop here. We have, in fact, barely scratched the surface. But, if you think that the ability to understand ambiguous sentences is a purely linguistic ability, think again. You can, for example, readily construe the cube to the left of Figure 1.1 from two perspectives: (a) and (b).¹⁶




Figure 1.1 Two visual interpretations of the same outline cube

And here too, juxtaposing two ambiguous objects appears to demand a parallel interpretation. You must see the cube on the left in Figure 1.2 from the same perspective you see the cube on the right.¹⁷ In the same vein, you may have felt in cases like (4) and (5) that the sentences contained redundant material, and you may have felt to desire to shorten them like this: *John can ride a bike, and Mary can, too; John can ride a bike, but Mary cannot.*



Figure 1.2 Viewing objects from the same perspective

This ability to understand material that is not explicitly present is also part of your tacit linguistic knowledge, but it shows up elsewhere in cognition. For example, everyone I have met interprets the visual object in Figure 1.3a as a white strip covering a black strip.¹⁸ You would, I'm sure, be surprised, if there was nothing underneath the white strip; that is, if you found what is in Figure 1.3b. You would be equally surprised if there was something as weird as what is in Figure 1.3c




Figure 1.3 What's underneath the white strip?

underneath the white strip. This shows that your visual system imposes severe constraints on object construal. This is true of your linguistic system too. You would be surprised if I told you that *John can ride a bike but Mary cannot* means something like “John can ride a bike, but Mary cannot drive a car.”

4 A Few More Notions

Linguists and other cognitive scientists are trying to figure out what these constraints on interpretation are, how exactly they work, and why they should be part of our mind. In so doing, linguists practice what is known as *descriptive grammar*, as opposed to *prescriptive grammar*.¹⁹ Prescriptive grammar is what people tend to associate with the term “grammar.” It refers to a normative practice that seeks to dictate proper linguistic usage. I guess you can say that this is the sort of thing that everyone ignores, unless they have to draft an official letter, make a public address, and so on. Prescriptive grammarians will tell you never to end sentences with a preposition, and you’d be right to wonder what those guys are talking about [sic]. Descriptive grammarians do not impose artificial rules that no one follows; they instead focus on normal usage, and try to figure out the underlying capacity that is responsible for what is said, and how it is interpreted. In so doing, descriptive grammarians also focus on what is sometimes called “negative” knowledge, or knowledge of negative facts: things that speakers never do. We already saw a negative fact in the context of the Parallelism Constraint (you *cannot* interpret sentences like *John can ride a bike and Mary can too* as four-way ambiguous). Here is another. Whereas native speakers of English commonly end their sentences with prepositions (*What did you say that Mary was talking about?*), they never say things like *What did you say about that Mary was talking?* For some reason, it is OK (cognitively/descriptively, but not prescriptively) to leave a preposition dangling at the end of a sentence, but it is not OK (cognitively/descriptively, and – irrelevantly – prescriptively) to leave a preposition dangling in the middle of a sentence. Prescriptive grammarians never worry about such negative facts (since no one ever does it, there is nothing to prohibit), but descriptive grammarians use them all the time to discover the limits of linguistic knowledge, the boundaries of our language faculty.

In some cases the limits will not be strictly linguistic, they will be shared with other cognitive abilities. This is arguably the case for the Parallelism Constraint. In such situations, the constraint will be said to be part of our language faculty broadly construed (*Faculty of Language Broad*, or *FLB*). But in some other cases, the negative knowledge will be strictly linguistic (part of our language faculty narrowly construed; *Faculty of Language Narrow*, or *FLN*).²⁰ (Teasing apart the contributions of FLN and FLB is one of the most exciting ways of figuring out how the mind works. For this reason I will devote a whole chapter (Chapter 12) to the distinction.)

One of the clearest examples of such an FLN-constraint comes from the way we interpret pronouns (elements like *he*, *she*, *it*, etc.). Native speakers of English have very sharp intuitions about sentences like (6–9).

- (6) John likes him.
- (7) He likes John.
- (8) John said that he likes Mary.
- (9) He said that John likes Mary.

In a normal context (with normal intonation), *John* and *him* cannot be understood as referring to the same person (meaning “John likes himself”). The same is true of *he* and *John* in (7). But in (8), *John* and *he* can “co-refer,” though not in the minimally different sentence in (9).

Sentences like (10) and (11) (where co-reference between *John* and *he* is possible in both sentences) make clear that whatever the principle at play is, it cannot be something very simple and intuitive. And yet, speakers’ intuitions are remarkably uniform, and consistent.

- (10) When John came in, he sat down.
- (11) When he came in, John sat down.

When analyzing sentences like (6–11), linguists invariably appeal to technical terms like dependent clause, embedded clause, and so on – constructs that I do not want to go into at this stage (I will come back to them in later chapters).²¹ The take-home message for now is that, as far as one can tell, there is no way to account for the paradigm in (6–11) without appealing to strictly grammatical notions – which speakers are not aware of, but which, at some level, must exist if we are to account for how they understand sentences. (Even something like “pronoun” is a concept or category that speakers must rely on to formulate the rule.)

5 Problems and Mysteries

From the discussion so far, it is clear that we are such thoroughly linguistic animals that we hardly realize what a complicated business language is. I readily confess that some of the sentences that linguists focus on are not very frequently uttered, but in so doing, linguists are no different from other scientists, who try to isolate hidden principles by performing artificial, and highly contrived experiments. Linguists are lucky that they can run interesting experiments very quickly, by probing speakers’

intuitions about sentences they may have never heard before. For here lies one of the most amazing facts about our linguistic knowledge: we are able to produce and understand sentences that we have never heard before. This is sometimes referred to as *the creative aspect of language use*.²² The most conspicuous manifestation of it is our ability to produce (and understand) ever-longer sentences:

- (12) John did it.
 Mary said that John did it.
 Bill claimed that Mary said that John did it.
 Sue believes that Bill claimed that Mary said that John did it.
 Harry thinks that Sue believed that Bill claimed that Mary said that John did it.
 ...

And speakers can readily judge sentences as acceptable, unacceptable, good, awkward, impossible, etc. no matter how unfamiliar the sentences are. As Chomsky famously noted, native speakers of English can judge a sentence like *Colorless green ideas sleep furiously* as a grammatically possible sentence of English (though it's hard to know what it could possibly mean), and can readily distinguish it from *Green sleep ideas furiously colorless*, which is not only meaningless, but also utterly impossible (linguists like to call such examples "word salad").

It is such intuitions (technically known as *acceptability judgments*) that linguists rely on in their attempt to uncover the principles responsible for mental structures. I think it is fair to say that few would have suspected the richness that 50 years of intensive research on mental structure have revealed. And I hope to reveal some of that richness here, and, hopefully, convince you that you too can contribute to this science.

Before concluding this chapter with an overview of the rest of the volume, I want to make a cautionary note. Although I hope that the present book will convince you of how much can be learned about human nature by studying human language, it is also important to remain aware of the fact that discovering mental faculties, such as the faculty of language, goes hand in hand with discovering their limits. Although the research program carried out by modern linguists has proven remarkably successful in some respects, it has also proven quite limited in other respects. I do not want to make false promises: This book will provide little, if any, insight into the nature of Consciousness, the Self, Meaning, and other topics that the layman tends to associate with "the Mind."²³

Cognitive science, like science in general, does not intend to be a theory of everything. The path of science is usually one of humility, a journey during which one learns to lower one's expectations, and appreciate how deep one can go into a few simple topics. The great physicist Richard Feynman was right when he said that the reason physics has been successful is because it focuses on a few simple things (leaving the rest to the chemists, who in turn pass it on to the biologists, who in turn have passed it on to the cognitive scientists, who in turn pass it on to

the novelists).²⁴ The Ancient Greeks thought that it would be easier to discover principles of the mind than principles of the heavens, because the mind was somehow “closer” and therefore more readily accessible to us. History has taught us that just the opposite is true. It may well be that, as Descartes suspects, we “do not have intelligence enough”²⁵ to understand some of the things that we are most curious about. The tale that follows emphasizes the existence and richness of cognitive abilities. In so doing, it also highlights their limitations. It may very well be that the very same cognitive abilities that enable us to understand some aspects of the world set us on a wrong path when it comes to other aspects of the world. As Noam Chomsky often remarks,²⁶ the very same organ (the brain) that defines what can be a problem for us (something which we can investigate with some hope of understanding) also defines what must remain a mystery (something which we cannot even turn into a problem). Just like it’s physically impossible for us to do certain things, there are things that are cognitively impossible for us to achieve; that means that there are some questions we will never answer.

6 Organization

With all these preliminary remarks in mind, let me now touch on the structure of the book. I have decided to organize the material around the five questions that have defined linguistic inquiry over the past 50 years. These are:²⁷

- 1 What is the best characterization of our Knowledge of Language?
- 2 How is that Knowledge acquired?
- 3 How is that Knowledge put to use?
- 4 How is that Knowledge implemented in the brain?
- 5 How did that Knowledge emerge in the species?

These five questions touch on many disciplines, beyond traditional areas of linguistics: evolutionary biology, neuroscience, developmental psychology, and more. Linguistic inquiry, guided by these questions, is interdisciplinary science in the strongest sense. It forces scientists to join forces, and to learn from each other.

Inquiry guided by the five questions just listed has made substantial progress. The questions have led to the best theory developed about language and its place in the human mind, and, as we will see in subsequent chapters, they can be taken as a model to study other cognitive abilities. Cognitive science is still very much in its infancy, and many cognitive abilities remain untouched. Recent work on music by Ray Jackendoff and Fred Lerdahl,²⁸ and on morality by John Mikhail, Sue Dwyer, and Marc Hauser,²⁹ have shown that adopting the methodology of linguistics (“pushing the linguistic analogy,” as it is sometimes described) can lead to important insights into the structure of other mental faculties.

At the same time, the answers provided by linguists have reached a certain depth that make them suitable for use as bridges to close the gap between our current understanding of the mind and of the brain. Although Spinoza in the seventeenth century already stated that the mind is the idea of the brain,³⁰ a formulation echoed by Steve Pinker's definition of the mind as what the brain does,³¹ we are in need of linking hypotheses connecting brain and mind. But as researchers like David Poeppel have stressed for years,³² such linking hypotheses are hard to come by because the concepts used in mind science and in brain science are not commensurable. It is hard for neurobiologists to know what to look for in the brain when what they are given by psychologists doesn't quite match what they already understand about the brain. One can only hope that as the tools and results of cognitive science are refined, they will become usable to formulate precise questions about the brain. There are signs, from very recent studies, which I will review in due time, that linguistics may be getting closer to defining an actual research program for neuroscience. It is still very preliminary, but I think there are grounds for optimism.

In sum, linguistics will play three roles in the following pages: it will be used as a theory (of a particular aspect of human cognition, the language faculty), as a model (to investigate other aspects of human cognition), and as a program (to formulate questions about how exactly the brain produces the mind). If I am correct about this, linguistics ought to be seen as one of the core areas of the Science of Man, an area that no serious cognitive scientist can afford to ignore.

To help the reader better appreciate where we are now, and what we can reasonably expect in the future, it is very useful to examine the context of emergence of modern cognitive science. For this reason the next chapter discusses the major factors that led to what is often called the "cognitive revolution"³³ (or "the mind's new science")³⁴ that took place in the 1950s. I then turn to a discussion of phenomena that taken as a whole provide some of the strongest reasons to investigate the biological foundations of language, which motivate modern linguistics.

The next two chapters thus alternate between the very general bird's-eye-view on cognitive science, and the more narrow, sometimes even microscopic, focus on the nature of the language faculty. In so doing I seek to establish a rhythm that I try hard to maintain in subsequent parts of the book, alternating between the laserbeam (focus on the details) and the floodlight (highlight the big picture). Although language will often be foregrounded, the reader should never lose sight of the fact that this is all in aid of a much bigger research project, one that promises to deliver insights into our most fundamental nature as a species. It's the bigger picture that Einstein urged us never to forget while doing our equations.

2

The Mechanization of the Mind Picture

Philosophers and historians are keen to point out that studying history may help one avoid repeating the mistakes of the past.¹ They are right. History is also a good indicator of the current state of the art (and the science); it's the best way to find out whether current hypotheses are as portentous and astonishing as we may think they are. In short, history and philosophy – two disciplines so often ignored by many scientists – are two of the most useful guides in the conduct of scientific inquiry. For this reason I would like to devote an entire chapter to discussing the major ingredients of the “cognitive revolution,” via which modern cognitive science established itself. Needless to say, a chapter won't be enough to touch on all aspects of the cognitive revolution. Thousands of pages could (and have) been filled with details of how modern cognitive science came to be.² So I will have to be highly selective, and extract some of the most important and (in the context of this book) most directly relevant lessons that can be gathered from this very important period in intellectual history.

1 Four Central Ingredients

I think it's fair to say that cognitive science as we know it today emerged in the mid-1950s at the confluence of:

- 1 the revival of and renewed appreciation for insights from what may well be called the “first cognitive revolution” that took place in the era of the enlightenment/modern period (seventeenth and eighteenth centuries);
- 2 the solidification of the scientific study of behavioral instincts in animals (ethology);
- 3 progress in the domain of mathematical understanding of notions like computation and information, algorithms, recursive functions, and the like; and

- 4 dissatisfaction with the then-dominant behaviorist paradigm in psychology, which took external behavior not as evidence for what goes on inside, but as the basic limit of inquiry.

Chomsky's early work contributed to all four strands of research that constitute the conceptual underpinnings of modern cognitive science.³ But he was by no means alone. And linguistics was certainly not the only field affected.

Today, rationalist, biological, computational, and mentalist considerations not only constitute the foundation of the "computational-representational"⁴ theory of mind, which many rightfully take to be the best hypothesis we have when it comes to finding out how the mind works. Together, they constitute a conceptual core that I would characterize as "non-negotiable."

I will touch on the four aspects listed in points 1–4 above, beginning with the reaction against behaviorism. I will then turn my attention to its immediate alternative, ethology, which allowed the mind and the mental to be regained. Next I show how extending the ethology program to human cognitive behavior in many respects revives concerns first articulated by giants like Descartes, Hobbes, and others, who sought to (in Hume's terms) "introduce the experimental method of reasoning into moral subjects"⁵ – that is, extend the revolution of the world picture brought about by Descartes, Galileo, and culminating with Newton, to the domain of the mind; in other words, develop a Science of Man. I conclude with a discussion of the mathematization of concepts like computation and information achieved by Alan Turing and Claude Shannon in the first part of the twentieth century, which laid the formal foundation for how to approach the mind in a rigorous (i.e., scientific) fashion.

I should point out that some of the themes I touch on are among the richest in intellectual history, and I cannot hope to be exhaustive even in those aspects of the cognitive revolution that I selected. My aim is very modest. I simply want to give the reader a sense of what it took to lay the foundations of modern cognitive studies.

2 The Demise of Behaviorism

In many ways, the cognitive revolution was a reaction to behaviorism, the then-dominant paradigm in psychological research. To the curious minds of young students like Chomsky, behaviorism was a bankrupt research program, as he made abundantly and unambiguously clear in his review of B. F. Skinner's book *Verbal Behavior*⁶ (Skinner was a major proponent of the behaviorist school, and his 1957 book was meant to be a synthesis of behaviorists' achievements in the domain of language). Chomsky's 1959 review of the book is uniformly seen as a classic document in the birth of modern cognitive science, and regarded by many as the nail that closed the coffin of behaviorism, and allowed the mental realm to be regained as an object of scientific inquiry.

There is a lot of truth to this picture, although I should point out that it is in part the result of parochialism. Behaviorism was overwhelmingly present in the geographical context in which modern cognitive science emerged (the US East Coast), but the situation in, say, Europe was very different, and in many ways, much more congenial and receptive to the mentalist position about to take over in the US.⁷

2.1 Behaviorism in a nutshell

Behaviorism purported to develop a scientific account of behavior and its acquisition.⁸ Its defining (and in many ways, most bizarre) characteristic was its self-imposed restriction to “observables.” According to major proponents of behaviorism such as B. F. Skinner, behavior was to be accounted for strictly in terms of stimuli coming from the environment, and the responses coming from the organism under study. The organism itself was treated like a blackbox, whose internal constitution (its internal structure) was deemed unfit for proper scientific inquiry. For behaviorism, the inner workings of the mind/brain were the philosopher’s fiction. Science had to proceed on the basis of a rigid scientific method that admitted of no “hidden,” “abstract” entities. The model to understand behavior was Pavlov, and his famous experiments with dogs.⁹ The scientist would note a correlation between a certain stimulus (the appearance of food) and a response from the organism (the dog salivating). The scientist would then introduce a new element into the environment: whenever food appears, a bell rings. Very quickly, the dog learns the connection between the appearance of food and the bell, and starts salivating as soon as the bell rings, even in the absence of food. The dog has thus learned to form a novel (arbitrary) association.

Behaviorists were fond of another experiment: Let the organism behave naturally at first (e.g., let a pigeon peck for food anywhere in the cage it finds itself). Select a pecking location at random, and lead the pigeon to focus on that particular location by rewarding its action when appropriate (pecking on the designated spot), and punishing its action when inappropriate (pecking elsewhere). The combination of reward and punishment will eventually result in the desired behavior becoming a learned habit.

Behaviorism thus proposed an extremely simple, hence at first appealing theory of learning, couched in terms like stimulus, response, and reinforcement (reward/encouragement and punishment/correction). The theory could generalize to any sort of behavior, and to any animal: whatever worked for teaching a dog to salivate at the right time, or a pigeon to peck at the right location, could be extended to teaching a child to give the right verbal response in the right context. The basic idea is that a child would acquire her vocabulary in the same way, for instance by learning to say “book” in the presence of a book-stimulus in the environment. The contribution of the inner nature of the organism could be ignored; the mind was not so much explained as explained away. Language learning (appropriate

verbal behavior) was a matter of associating the right sounds/words with the right objects out there in the world (the stimuli), through the judicious intervention of “teachers” correcting or rewarding responses.

The task of the scientist (the linguist, if verbal behavior was at issue) consisted solely of identifying the right stimuli, giving rise to the right responses. The learning mechanism was maximally general: neither domain-specific nor species-specific.

2.2 Chomsky’s review

In his review of Skinner’s book, Chomsky took issue with each and every one of Skinner’s claims. Chomsky’s points turned on several themes that have since characterized his approach to language:

- 1 Studied in behaviorist terms, the study of language reveals nothing about the mind (the internal structure of the species), as it concentrates exclusively on the environment. Nothing is said about the inborn structure that must surely be present to account for the very possibility of behavior.
- 2 Unlike other sciences, which do not predefine the terms in which all explanation must be couched, behaviorism imposes severe limits both in terms of what can be studied, and how it should be explained.
- 3 Learning is a much more subtle and complex business than a mere matter of reinforcement (reward/punishment); the maturation of the organism, and the lack and/or inefficiency of correction (i.e., the absence of explicit teaching) must be acknowledged (think of the examples of knowledge “unlearned and untaught” mentioned in Chapter 1).
- 4 Above all, in the domain of language, the creative aspect of language use must occupy center stage. When presented with a given stimulus (say, a Rembrandt painting, to use Skinner’s own example), a speaker is not compelled to say “Dutch” (as Skinner claimed), but can offer an infinity of responses, most of which are seemingly unrelated to the stimulus (say, “I’m hungry”). The speaker even has the choice of saying nothing.
- 5 In connection with the creativity argument, Chomsky points out that linguistic behavior is radically independent of any notion of probability. Since the response is not determined by the stimulus, the frequency of a given stimulus in the environment cannot lead to a good prediction of the response. Speakers are not machines, but endless sources of novel utterances.
- 6 The ability to make linguistic sense of stimuli, to distinguish “news” from “noise,” amounts to handling the stimulus (information) in highly specific ways, which a domain-general, and species-general theory of learning cannot even begin to account for. The organism must somehow be equipped with an innate disposition to learn from the environment, to pick out (independently of reinforcement) the relevant aspects of the stimuli. It won’t do to treat the organism as a blank slate.

- 7 Verbal behavior (the “response”) is highly structured (e.g., words in sentences are not simply strung together like beads on a string (more on this in Part II)), in a way that is not reflected at all in the environment. In other words, the organism structures the input (stimulus) according to its own rules (innate predispositions).
- 8 The behaviorist account of behavior fails to do justice to the extraordinary richness of the “response.”

Chomsky’s review of Skinner’s *Verbal Behavior* is one of his most clearly articulated statements on the nature of language, and what linguistics should aim at. Together with the introductory chapter of his 1965 book *Aspects of the Theory of Syntax*, I think it should be on all the required-readings lists for courses on language and the mind. But too often, the review is narrowly characterized as a compelling critique of behaviorism – which it is, but this captures only half of the essence of the review. Chomsky’s essay is far more than a negative piece. In addition to revealing the lethal limitations of behaviorism as a science of behavior and learning, Chomsky sketches the foundation of what would be needed to account for the central aspects of language (and cognition). Throughout, Chomsky stresses the need to adopt a mentalist and nativist stance, to reveal the built-in structure in terms of which the organism processes information in highly specific ways. In so doing he cites approvingly the work of ethologists (to be discussed below) like Niko Tinbergen, Konrad Lorenz, and W. H. Thorpe, who studied complex innate behavior patterns in animals, and their “innate tendencies to learn in specific ways.” Chomsky insists that the development of our cognitive faculties isn’t to be explained in terms of learning; rather, such faculties grow, like other organs of the body. Of course, the environment is needed, just like proper nutrition is required for the body to grow. But the internal engine (our biological endowment) is the main source of the characteristics of our behavior. The following example might help here: We know that the beautiful pink color displayed by the flamingo is the result of the combination of shrimp and plankton in its diet. That’s, if you wish, the environment’s contribution to the end product, but there must be something more, something internal, going on, for no seagull fed shrimp and plankton would ever turn pink. The same is true in the realm of cognition. It’s the combination of internal and external factors that together can account for the end product. It was the behaviorists’ mistake to dispense with the internal entirely.

Chomsky also points out that the very fact that we do not yet have a satisfactory picture of the precise underlying biological and neural mechanisms underlying behavior (this was true in 1959, and it remains true today) should not prevent us from positing mental entities with which we can shed light on our human nature. After all, Newton’s inability to understand the true nature of the gravitational constant didn’t prevent him from positing g in his equations and thereby providing us with a deeper understanding of planetary motion, tides, and much more.¹⁰ It takes a good scientist to know when is a good time to say *hypotheses non fingo*.


Finally, Chomsky also stresses the importance of going beyond the surface appearance of the response, and studying its internal properties. In so doing, he cites Lashley's work¹¹ on the problem of "serial order" in behavior. Lashley, once a behaviorist, was at the time working on showing its limits, by emphasizing the intricacy of the internal structure of the "response," the syntax (in a broad sense) of behavior. By pointing to Lashley's work, Chomsky (who himself had written about the internal constituency of language in his massive 1955 *Logical Structure of Linguistic Theory*,¹² and in *Syntactic Structures*¹³) was drawing attention to the need to first figure out the nature of knowledge before tackling the question of acquisition of that knowledge. This too has remained a theme in Chomsky's writings to the present day.¹⁴

In sum, Chomsky is thus not merely rejecting the behaviorist program, he is urging us to adopt an alternative program that does justice to the mental, to the inner structures of the organism, and to the exquisite (invisible) structure of the response. Chomsky's 1959 review of Skinner's book was his first explicit (published) statement of a mentalist/nativist position from which he has never deviated. His subsequent work, and that of many associates, can be seen as one long argument for, and refinement of, this position, with a few shifts of emphasis over the years, but overall remarkably consistent in its ultimate goals.¹⁵ For example, we'll see in subsequent chapters how Chomsky has remained highly skeptical and critical of appeals to properties of the environment to explain properties of language and the mind (this will be the case in the domain of "linguistic meaning" (Part III), and in the domain of "language evolution" (Part IV).)

To the modern mind, it is very hard to believe how behaviorism could have been so important, so dominant. Whenever I cover behaviorism in one of my classes, exposing all its limitations, my students ask me who in their right mind would think that something as simplistic as a stimulus–response–reinforcement model would be adequate. And yet, behaviorist tendencies still rule in many corners of cognitive science. In a more sophisticated form, to be sure, but the appeal to common-sense notions like explicit instruction and hard-won learning from the environment often proves too strong. A quick look at Steven Pinker's book *The Blank Slate*¹⁶ reveals how widespread behaviorist/empiricist assumptions continue to be in the study of human nature.¹⁷ Perhaps, as developmental psychologist Lila Gleitman shrewdly noted,¹⁸ empiricism (which banishes talk of the innate) is innate. It is also wrong, but it may be one of these biases of our nature that invariably tempt us to adopt the path of least resistance, and assume that our knowledge is all too easily explainable in terms of teaching and learning. Paul Bloom may be right when he says¹⁹ that the very kind of mind we have may be the source of our recurrent mistake of treating the mind as something special, different from the body, which no one takes to be a *tabula rasa*. (Would anyone suggest that seeing, walking, digesting, and the like are learned faculties?)

The point bears emphasis. Modern cognitive science is, I think, on the right track. But it remains an experiment: one which rests on a rock-solid basis of conceptual and empirical arguments, but which nevertheless remains in a precarious

state, as it relies on the exercise of an organ whose internal constitution all too often leads to the illusion of common-sense wisdom that all it takes is a good teacher and hard work. This would be one of those inevitable illusions that Massimo Piattelli-Palmarini has so eloquently discussed,²⁰ on a par with our irresistible urge to interpret the horizontal line in (1) as being longer than the (in actuality, exact same) line in (2) (the so-called Lyell-Müller illusion).²¹



3 The Ethologist Program

Whereas Skinner took inspiration from Pavlov's experiments with animals to account for verbal behavior, Chomsky took inspiration from the ethologists, who studied and rationalized animal behavior in a radically different, and (from our perspective) much more fruitful perspective.


3.1 The contribution of the organism

Chomsky's 1959 review, as well as the work of his close associate Eric Lenneberg, the author of the classic *Biological Foundations of Language*,²² is peppered with "ethologist talk." Ethologists maintained that animals possess predictable behavioral programs; they have a basic repertoire of behavior patterns, which mature during the course of development. The emergence of these movement patterns, like that of cells and organs, is guided by phylogenetically acquired blueprints.

The subject matter of ethology was first presented in textbook form by Tinbergen in 1951,²³ and couldn't have been more different from the behaviorist paradigm. As we saw, behaviorism essentially amounted to doing psychology without the mind; seeking to explain all behavior on the basis of reflexes, repudiating terms like feeling, attention, and will, and asserting instead that one can only determine stimuli and reactions and the laws governing their interactions. By stressing the influence of the environment, behaviorists overlooked the inherited, innate basis of behavior. In so doing, they made the fatal mistake of elevating one explanatory principle to an all-exclusive one.

For a long time, behaviorists overlooked the spontaneity of behavior, not readily observable in their particular experimental settings; for them all behavior consisted of reactions to stimulation. Ethologists proved this research strategy to be much too one-sided. They showed in unambiguous ways that behavior could not

be assumed to be merely a response to external stimuli. In between the Stimulus and the Response, there must lie some innate faculty (which, in the context of language, Chomsky called the “Language Acquisition Device” (LAD)²⁴ or “Universal Grammar” (UG)²⁵).



It was because they took the basis of behavior to be innate (it was, for them, a biological *a priori*), that ethologists focused on the study of instincts (which William James defined as the “correlates of organs/structures”²⁶).

They showed,²⁷ for example, that many young animals exhibit an innate avoidance of a precipice, which they recognize visually before having had the adverse experience of falling off a cliff. (Chomsky makes a similar point in his review of Skinner’s book, when he notes that humans can respond to the verbal stimulus “give me your money or I’ll kill you,” even if they could not have experienced death beforehand.²⁸)

Among the most famous experiments from the ethology literature is Lorenz’s study of imprinting.²⁹ Lorenz showed that the combination of a rhythmic call and virtually any moving object releases the following reaction in a young graylag gosling shortly after hatching (also true of chicks and ducklings): it follows a man as readily as it would a goose or a moving box, and remains with it. It cannot be induced to follow even its own mother once it gets imprinted to the man or the moving box.

It is as if the gosling is born with the knowledge “assume that the first thing that moves is your mother, and follow it.” Once acquired, the knowledge is retained for life (in contrast to prototypical instances of learning, where forgetting is common). What Lorenz discovered is an inborn disposition to learn, and therefore, the need to tease apart what he called “instinct–learning intercalation” (which he did by running “deprivation experiments,” meant to measure the importance of the environmental stimulus).³⁰

It should be emphasized that when ethologists claim that the development of a particular behavior does not require certain experiential influences, they certainly do not mean that no experience at all is necessary. Sometimes, even self-stimulation is required³¹ (as appears to be the case with language, where babbling is a necessary developmental stage in the normal individual, as we will see later on in this book; the same is true for birds when they learn to sing).³²

Ethologists had a very sophisticated view on the role of the environment. The great ethologists like Lorenz and Tinbergen were here following in the footsteps of Jakob von Uexküll,³³ whose experiments in the early twentieth century showed that an animal can perceive only a limited part of its potential environment with its sense organs; some of these perceived characteristics of the environment serve

as specific cues, or key stimuli, that match (like a lock and a key) the innate biases of the organism.

In other words, organisms come equipped with data-processing mechanisms and detectors that are fine-tuned to specific environmental stimulus situations (as Uexküll notes, animals construct their environment, their *Umwelt*). These highly specialized data-processing mechanisms enable the organism to react in a biologically adequate way upon the first encounter with the stimulus. For example, Lorenz observed³⁴ that a toy model (mimicking an approaching flying creature) used in an experiment with young birds in a nest reveals that the form of the object is largely irrelevant, but the relative speed with which the object approaches matters. In other words, the young birds are responding only to a certain kind of stimulus configuration.

The capacity to react to often highly specific stimuli prior to any individual experience with specific behavior patterns has been documented in numerous contexts by the ethologists. Chaffinches given a choice prefer the song of their own species, based on innate knowledge as to the type of song to be learned. Chimpanzees know innately that threats are made by producing noise (although they must learn the method of producing it). And a zebra finch learns its song from those who feed it. If a society finch feeds zebra finch young, they will learn the society finch song, even if there are zebra finches singing in an adjacent cage; however, if the young are fed by both species, they will learn the zebra finch song, so they still show a preference for their own species song, despite their flexibility.³⁵ (Perhaps this preference is nature's way of making sure that "wrong" imprinting is not too frequent.)

Another sign that animals clearly "construct" their environment (make their *Umwelt* out of the *Welt* out there, as it were) is that an animal does not indiscriminately associate each environmental stimulus with specific perceptions. For example, a rat that has tasted poisoned bait will avoid the bait, but not the place where it was found.³⁶ The ability to identify (and ignore) "irrelevant" aspects of the environment is perhaps one of the most underestimated features of the innate contribution of the animal to the learning task.

Ethologists discovered that not only do animals respond to specific stimuli, they possess their specific learning capacities only for a very limited time, after which the animal can no longer learn. In other words, the capacity to learn is associated with a detailed learning schedule. (We'll see in the next chapter that this is true of human language too, as Lenneberg was the first to suggest.)³⁷ In many publications ethologists suggested extending their study of animal behavior to human behavior, but it took a while until human ethology emerged as a scientific discipline.³⁸

For a long time the science of man has been dominated by the environmentalist assumptions, according to which all human behavior with the exception of some basic reflexes is learned, but animal behavior shows that behavior is preprogrammed in well-defined ways. From a modern, Darwinian perspective, we should expect that, if anything, humans have more instincts, more fixed action patterns than "lower" animals. That so many continue to assume that humans are blank slates, that because they are flexible they must be malleable, means that we haven't yet

assimilated Darwin's great message that we are no different from other creatures. We too are the products of our biology.

Darwin (in many ways a forefather of ethology, especially in his emphasis on the comparative method in studying behavior) already noted in *The Descent of Man*³⁹ that "man has an instinctive tendency to speak, as we see in the babble of our young children; whilst no child has an instinctive tendency to brew, bake, and write." (Darwin goes on to draw an illuminating analogy between human language, in all its diversity, and birdsong.)

3.2 The basic research questions

In his classic paper on "Aims and Methods of Ethology,"⁴⁰ Tinbergen organizes the research strategy of ethology around four questions:

- 1 What stimulates the animal to respond with the behavior it displays, and what are the response mechanisms?
- 2 How does an organism develop as the individual matures?
- 3 Why is the behavior necessary for the animal's success and how does evolution act on that behavior?
- 4 How has a particular behavior evolved through time? Can we trace a common behavior of two species back to their common ancestor?

These four questions correspond point by point to those that define the research program of modern linguistics (stated at the end of the last chapter):

- 1 What is the best characterization of our Knowledge of Language?
- 2 How is that Knowledge acquired?
- 3 How is that Knowledge put to use?
- 4 How is that Knowledge implemented in the brain?
- 5 How did that Knowledge emerge in the species?

These five questions were first stated in this very format in Chomsky's work of the mid-1980s,⁴¹ but his desire to carry out the ethologist's program in the realm of cognition can be traced back to his earliest writings. It is because of Chomsky's vision that we can use phrases like "the language instinct"⁴² or "the language organ"⁴³ as titles of books to refer to the subject matter of linguistics (and, similarly, for other domains of cognition; cf. Hauser's use of the term "moral organ"⁴⁴ in his study of our sense of right and wrong).

In his discussion of "learning organs" (including "the language-learning organ"), Randy Gallistel⁴⁵ provides a useful analogy that I cannot resist reproducing. He notes that Harvey revolutionized physiological thinking in 1628 when he showed that the heart circulates the blood and that its structure makes it possible for it to perform this function. In Gallistel's words,

Before Harvey, the modern conception of an organ as something whose particular structure enables it to perform a particular function did not exist. Physiological thinking centered not on organs but on humors. Pathological conditions were thought to arise from an excess or deficiency in one or more of them [Hamlet suffered from an excess of black bile, which was the cause of his melancholy]. Crucially, humors did not have specific and limited functions. Much less did they have a structure that enabled them to perform a specified function. Organs, by contrast, have specific and distinct functions and a structure that enables them to perform them.

Gallistel is right to say that when, building on the work of the ethologists, Chomsky reconceptualized language learning in terms of the language organ, his reconceptualization was as radical in its implications for psychology and neuroscience as Harvey's work was for physiology (a reconceptualization that has yet to be fully appreciated, as we will see in Part IV).

People generally conceive of learning (especially in humans) as mediated by a small number of very general learning processes, such as analogy, imitation, association – none of them tailored to the learning of a particular kind of material. These are so-called domain-general theories of learning, of the sort advocated by the behaviorists. Like the medieval humors, these learning processes do not have specific functions, nor do they have structures that enable them to perform those functions. In this view of learning, the brain is plastic, so much so that it can, for all intents and purposes, be equated with a blank slate. It rewires itself to adapt and adjust to experience. There are no problem-specific learning organs, computing representations of different aspects of the world from different aspects of the animal's experience.

Chomsky, by contrast, suggested that learning is mediated by distinct learning organs, each with a structure that enables it to learn a particular kind of contingent fact about the world; each with the ability to construct a specific *Umwelt*. He further claimed that much of what we know is not learned; rather, it is implicit in the structure of the organs that learn, which is what makes it possible for each such organ to learn in the first place.

As Chomsky has pointed out on numerous occasions,⁴⁶ whether this organ resides in a highly localized part of the brain or arises from an interconnection of diverse specific domains distributed all over the brain is irrelevant to whether it constitutes a distinct organ or not. Some organs are localized (for example, the kidney) while others are not (for example, the circulatory system). As Gallistel emphasizes, the essential feature of an organ is that it has a function distinct from the function of other organs and a structure suited to that function, a structure that makes it possible for it to do its job.

Although Chomsky had language foremost in mind, he clearly understood that his conception of learning implied that other forms of learning must likewise be mediated by problem-specific learning organs, and it is a tribute to him that in recent years learning organs have been posited to account for basic arithmetic, music, morality, and (I expect) many more (see Chapter 8).⁴⁷

Gallistel further observes that:

Chomsky understood that the notion of a general purpose learning process makes no more sense than the notion of a general purpose sensing organ – a bump in the middle of the forehead whose function is to sense things. There is no such bump, because picking up information from different kinds of stimuli – light, sound, chemical, mechanical, and so on – requires organs with structures shaped by the specific properties of the stimuli they process. The structure of an eye – including the neural circuitry in the retina and beyond – reflects in exquisite detail the laws of optics and the exigencies of extracting information about the world from reflected light. The same is true for the ear, where the exigencies of extracting information from emitted sounds dictates the many distinctive features of auditory organs. We see with eyes and hear with ears – rather than sensing through a general purpose sense organ – because sensing requires organs with modality-specific structure.

And I hope the reader never loses sight of the fact that even the simplest task of perception (say, visual perception) is a fearsomely complex task, based on fragmentary information detected by the retina (so fearsomely complex that to this day it remains far from well understood how the brain comes up with a unified perception of stable objects situated in a spatial environment, such as a book on a table).

Chomsky took the central message of the ethologists to be that learning different things about the world from different kinds of experience requires mental computations tailored both to what is to be learned and to the kind of experience from which it is to be learned. Therefore, there must be task-specific learning organs, with structures tailored both to what they are supposed to extract from experience and to the kind of experience from which they are to extract it. And it is the task of the cognitive scientist to figure out what the specific structures of these learning organs are. (As we will see in the next chapter, the built-in structure of the learning organ, which is what makes learning possible, has as a consequence that the inferences the learner draws from limited experience go beyond what is justified by that experience.)

Gallistel is again perfectly right when he notes that, for computational reasons, learning organs may be expected to differ between species of animals, just as do sensory organs. He notes that pit vipers sense infrared radiation, whereas we do not, because they have a sensory organ, which we lack. We learn languages, whereas pet fish do not, because we have a language-learning organ, which they lack. Gallistel has documented in his own work remarkable learning abilities in animals.⁴⁸ Even insect brains compute representations of behaviorally relevant aspects of the world. For example, they compute the animal's position in the world by integrating its velocity with respect to time. Other examples are the learning of the solar ephemeris, the construction of a cognitive map, and episodic memory in food caching. These are truly remarkable computational feats (many of which lie beyond our reach), which a program like behaviorism (or any other approach based on general learning) would invariably obscure.

4 Insights from the Cartesians

Although Chomsky took direct inspiration from the ethologists when it came to formulating a concrete research program, with experiments that would reveal the inner structure of learning organs, the knowledge that some behavior is relatively uninfluenced by individual experience is much older than the field of ethology.

In his brief but excellent historical introduction to ethology,⁴⁹ Eibl-Eibesfeldt reports that Baron von Pernau knew that animals possessed innate skills in addition to those they acquired, behavior patterns that they did not have to learn by imitation or other forms of training. In a work dated 1716, he described the behavior patterns of various birds and showed which species had to learn their songs from their parents and which, upon becoming sexually mature, were able to sing their species-typical songs without prior exposure to them.⁵⁰ Reimarus wrote in a similar vein in 1762:⁵¹

How do the spider and the ant lion go about finding means of supporting themselves? Both can do no other than to live by catching flying and creeping insects, although they are slower in their own movements than is the prey they seek out. But the former perceives within herself the ability and the drive to weave an artful net, before she had as much as seen or tasted a gnat, fly, or bee; and when one is caught in her net she knows how to secure and devour it . . . the ant lion on the other hand, who can hardly move in the dry sand, mines a hollow funnel by burrowing backward, in expectation of ants and worms that tumble down, or it buries them with a rain of sand that it throws up in order to cover them and bring them into his reach. . . . Since these animals possess by nature such skills in their voluntary actions that serve the preservation of themselves and their kind, and that admit many variations so they possess by nature certain innate skills . . . a great number of their artistic drives are performed at birth without error and without external experience, education or example and are thus inborn naturally and inherited . . . one part of these artistic drives is not expressed until a certain age and condition has been reached, or is performed only once in a lifetime, but even then it is done by all in a similar manner and with complete regularity. For these reasons these skills are not acquired by practice . . . but not everything is determined completely by the drives of the animals and frequently they adjust, of their own volition, their actions to meet various circumstances in various and extraordinary ways . . . for if everything and all of their natural powers were to be determined completely, that is, would possess the highest degree of determination, they would be lifeless and mechanical rather than endowed with the powers of living animals.

Chomsky again realized that wonderful insights like those expressed in this passage, so relevant to the concerns of modern cognitive science, would be lost if the anti-historical stance so common to the sciences were adopted. Chomsky was among the first to trace back current issues to the work of the natural philosophers of the early modern period; he did this specifically in the domain of linguistics (see his *Cartesian Linguistics*), but it is interesting to note reference to Descartes,

Spinoza, Hume, and others in the titles of many recent publications in cognitive science.⁵² (The title of this chapter too is an allusion to this fascinating period of history, well narrated in Dijksterhuis's *The Mechanization of the World Picture*.⁵³)

There is indeed a lot to learn from the rationalists, the empiricists, and the romantics in the domain of language in particular, and in the domain of cognition as a whole. Many of the debates and controversies raging today in cognitive science can be said to have been fought first in the seventeenth and eighteenth century – with different means, and slightly different arguments, to be sure, but with the same guiding intuitions, the same basic hypotheses, the same passion, and the same intention to shed light on the nature of human beings.

Here more than elsewhere I will have to confine my presentation to a few key aspects of the period. I will begin with Chomsky's seminal discussion of Descartes' legacy, then turn my attention to other key thinkers.⁵⁴

4.1 Descartes and (innate) representations

More than anyone else, Descartes emphasized the role of language in cognition – the study of language went along with a specific theory of mind, for, as we already saw, language was, in the words of Leibniz, “the best mirror of the human mind.” For him, the way we use language was the most obvious reflex of what is distinctively human in cognition. Descartes was quick to point out that the key difference between humans and other species was more than skin-deep, and stood in the way of reducing man to a machine. As he notes:⁵⁵

It is quite easy to conceive of a machine so constructed so that it utters words, and even words which correspond to bodily actions causing a change in its organs (for instance, if you touch it in one place it asks what you want of it; if you touch it in another place it cries out that you are hurting it, and so on). But it is not conceivable that such a machine would produce different arrangements of words so as to give an appropriately meaningful answer to whatever is said in its presence, as the dumbest of men can do.

Descartes stressed that the ability to use language must not be confused with “natural movements which express passions and which can be imitated by machines as well as by animals.” The crucial difference is that automata “could never use words or put together other signs as we do in order to declare our thoughts to others.” That capacity to express thoughts as we do is a distinctively human one, independent of intelligence:⁵⁶

It is quite remarkable that there are no men so dull-witted or stupid – and this includes even madmen – that they are incapable of arranging various words together and forming an utterance from them in order to make their thoughts understood; whereas there is no other animal, however perfect and well-endowed it may be, that can do the same.

Elsewhere, Descartes notes that this limitation on the part of animals:⁵⁷

does not happen because they lack the necessary organs, for we see that magpies and parrots can utter words as we do, and yet they cannot speak as we do: that is, they cannot show that they are thinking what they are saying. On the other hand, men born deaf and dumb, and thus deprived of speech-organs as much as the beasts, or even more so, normally invent their own signs to make themselves understood by those who, being regularly in their company, have the time to learn their language.

As Chomsky sums it up, “man has a species-specific capacity, a unique type of intellectual [cognitive] organization which cannot be attributed to peripheral organs or related to general intelligence, and which manifests itself in what we may refer to as the ‘creative aspect’ of ordinary language use – its property being both unbounded in scope and stimulus-free.”⁵⁸ As we saw already in the context of behaviorism, it is the detachment of the use of language from the immediate circumstances in which that capacity is exercised that Chomsky, with Descartes, sees as the essence of language and human cognition. Animals may use specific calls in the presence of predators, but only humans can form novel utterances involving unicorns, worlds in 11 dimensions, predictions for the twenty-ninth century, and reflections on our distant past.

It is interesting to note that Descartes was writing in an intellectual context not unlike the one in which modern cognitive science merged. Much as many people were (and still are) fascinated by computers and the possibility of artificial intelligence in the 1950s, many Europeans in Descartes’ time were enamored with automata. Celebrated engineers like Vaucansson built mechanical ducks that could digest, human-like machines that could dance and utter sounds when certain buttons were pressed, etc.⁵⁹ These machines were meant to shed light on the properties they were modeling, much like today’s computer simulations (think of the metaphor of the universe as clockwork). They were – literally – toy models, crucial to the development of modern science. Indeed, people like Descartes argued that mechanical explanations should be the aim of scientific understanding: if you can’t think of a way a machine could be built to give rise to whatever property you are studying, you can’t claim to have understood that property. It is for this reason that human language plays such a key role in Descartes’ philosophy, for in the creative aspect of language use Descartes found the single most powerful example of the limit of mechanical explanation. Not surprisingly, Descartes saw language as a litmus test to be used to draw a distinction between man and machine (an idea revived almost three centuries later by Alan Turing with his “imitation game” later known as the “Turing test”).⁶⁰

Following Descartes many philosophers discussed the constitutive function of language with respect to thought – the major theme of what Chomsky called Cartesian linguistics. This was true not only of followers of Descartes, who shared his rationalist persuasions (in particular, his nativist stance), but it was also true of opponents, empiricists like Hobbes,⁶¹ who agreed that language played a critical role in human cognition.

It may be useful to mention at this point that at least in the context of what Brook called the prehistory of cognitive science,⁶² the traditional empiricist–rationalist dichotomy is less useful than it is in the context of the debate between modern cognitivists and behaviorists. Classical empiricists like Hobbes and Hume (but also Berkeley and Locke) had a good deal more to say about the mind than the simple label of empiricism may lead us to think.⁶³ True, there was a basic opposition⁶⁴ between those who made a clear distinction between perception and cognition and claimed that concepts have structures that are not given immediately in the perceptual array (and are unlearned, i.e., “innate”), and those who ran together perception and cognition, claiming that concepts are features abstracted from the perceptual array. But the great empiricists, like Hume, recognized that “we must discover those parts of knowledge that are derived by the original hand of nature.”⁶⁵ In other words, Hume recognized that no system to which purpose can be assigned can be entirely without native dispositions; in their absence, the system would have no wedge into any action which merits the appellation of “behavior.” Jerry Fodor⁶⁶ even goes so far as to talk about Hume’s “Cartesian naturalism,” and urges us to go beyond the narrow characterization of Hume’s empiricist inclinations and appreciate the insights he has to offer on the architecture of our cognition once detached from his specific claims about where our concepts come from. I happen to think that Fodor is exactly right on this way of approaching Hume (and I would say the same about Hobbes); their works belong to Cartesian cognitive science, and have far more to contribute to it than the behaviorists.

Both rationalists and empiricists of the seventeenth and eighteenth centuries read a lot like modern developmental psychologists and cognitive scientists, at least in the questions they raise (if not in the answers they provide).

In Descartes’ wake, philosophers/grammarians took the study of language to be a branch of psychology and investigated ways in which the structure of language could reveal hidden layers bearing on the nature of cognition. Here the works of Port-Royal grammarians like Arnaud, Lancelot, du Marsais, and Beauzée deserve special mention, as do the works of Leibniz, Humboldt, and the German romantics like Schlegel. As Chomsky noted in *Cartesian Linguistics*, many of the technical details bear direct resemblance to some of the first analyses in modern linguistics following the cognitive revolution of the 1950s.⁶⁷

In addition to working on specific languages, these grammarians/philosophers also sought to determine the universal form of language, those principles that would be part of all languages. This search for universal principles and rational explanation went under the name of General or Universal Grammar, and the latter term has been incorporated into modern linguistics.⁶⁸

Along with its emphasis on the creative aspect of language use, Cartesian linguistics distinguishes itself from other approaches to language in its focus on mental structures (sometimes called “representations,” a term that is to be used with caution because it means different things to philosophers, psychologists, etc.). The Cartesians insisted that the sort of creativity and flexibility deployed in the context of language use (its detachment from immediate circumstances) required that humans be endowed

with innate concepts and cognitive faculties. In this way Cartesian linguistics becomes part of a more general concern of rationalist psychology with the contribution of internal (mental) factors to human knowledge and cognition.

As Herbert of Cherbury had noted as early as 1624, we bring to the world “principles or notions implanted in the mind” that are “stimulated by [external] objects,” but that no one is tempted to say are conveyed by the objects themselves.⁶⁹ It is these mental constructs that make experience and learning possible in the first place. In Herbert’s terms, these interpretive principles⁷⁰

are so far from being drawn from experience or observation that, without several of them, or at least one of them, we could have no experience at all nor be capable of observations. For if it had not been written in our soul that we should examine into the nature of things (and we do not derive this command from objects), and if we had not been endowed with Common Notions, to that end, we should never come to distinguish between things, or to grasp any general nature. Vacant forms, prodigies, and fearful images would pass meaninglessly and even dangerously before our minds, unless there existed within us, in the shape of notions imprinted in the mind, that analogous faculty by which we distinguish good from evil. From where else could we have received knowledge? In consequence, anyone who considers to what extent objects in their external relationship contribute to their correct perception; who seeks to estimate what is contributed by us, or to discover what is due to alien or accidental sources, or again to innate influences, or to factors arising from nature, will be led to refer to these principles. We listen to the voice of nature not only in our choice between what is good and evil, beneficial and harmful, but also in that external correspondence by which we distinguish truth from falsehood, we possess hidden faculties which when stimulated by objects quickly respond to them.

In focusing attention on the innate interpretive principles that are a precondition for experience and knowledge and in emphasizing that these are implicit (unconscious, tacit) and may require external stimulation to become active or available to introspection, Herbert formulated the core of the psychological theory in the immediate background of Cartesian linguistics, as pursued by Descartes, Leibniz, Spinoza, Cordemoy, the English Platonists like Cudworth, the German romantics (Schlegel, Herder), and Kant. All of them emphasized the need for the mind to structure experience, and were highly suspicious of the “senses.” As Descartes emphasized in his *Meditations*,⁷¹ sense perception is sense deception (think of all the instances of visual illusion). Descartes’ view is well articulated in this oft-cited passage:⁷²

. . . if we bear well in mind the scope of our senses and what it is exactly that reaches our faculty of thinking by way of them, we must admit that in no case are the ideas of things presented to us by the senses just as we form them in our thinking. So much so that there is nothing in our ideas which is not innate to the mind or the faculty of thinking, with the sole exception of those circumstances which relate to experience, such as the fact that we judge this or that idea which we now have immediately before our mind refers to a certain thing situated outside us. We make such a judgment

not because these things transmit the ideas to our mind through the sense organs, but because they transmit something which, at exactly that moment, gives the mind occasion to form these ideas by means of the faculty innate to it. Nothing reaches our mind from external objects through the sense organs except certain corporeal motions . . . But neither the motions themselves nor the figures arising from them are conceived by us exactly as they occur in the sense organs . . . Hence it follows that the very ideas of the motions themselves and of the figures are innate in us. The ideas of pain, colors, sounds, and the like must be all the more innate if, on the occasion of certain corporeal motions, our mind is capable of representing them to itself, for there is no similarity between these ideas and the corporeal motions. Is it possible to imagine anything more absurd than that all the common notions within our mind arise from such motions and cannot exist without them? I would like our author to tell me what the corporeal motion is that is capable of forming some common notion to the effect that “things which are equal to a third thing are equal to each other,” or any other he cares to take. For all such motions are particular, whereas the common notions are universal and bear no affinity with, or relation to, the motions.

In a similar vein, the Cambridge Platonist Ralph Cudworth⁷³ speaks of those “cognoscitive powers” that enable men to understand and judge what is received from the senses, and anticipate experience. In so doing, the rationalist psychologists were reviving insights of Plato’s,⁷⁴ “purging” (the verb used by Leibniz)⁷⁵ them from his dubious theory of pre-existence (the idea that innate ideas were gathered in previous lives), and opening the way to a more genetic source (which Darwin made clear in his famous statement: “Plato says in *Phaedo* that our necessary ideas arise from the preexistence of the soul, and are not derivable from experience – read monkeys for preexistence!”).⁷⁶

As Chomsky observes at the end of *Cartesian Linguistics*,⁷⁷ contemporary research in perception has returned to the investigation of the role of internally represented schemata or models, and has begun to elaborate the idea that it is not merely a store of schemata that function in perception but rather a system of fixed rules for generating such schemata. In this respect too, such recent works can be said to continue the Cartesian tradition in psychology, of which Cartesian linguistics is an essential part.

4.2 Hobbes and mental computations

Before leaving the prehistory of modern cognitive science there is one more idea that I’d like to touch on, and that is Hobbes’ claim that thinking/reasoning is computing, and more specifically, his claim that the very fabric of language is what supports this computation. Where Descartes stressed the importance of internal structures, Hobbes emphasized the idea that thinking amounted to performing arithmetic-like operations on these structures.⁷⁸ In so doing, Hobbes was advocating what came to be known in the past 50 years as the computational theory of mind, refined with the advent of precise notions of algorithms (see below).⁷⁹

Much like Descartes, Hobbes placed language at the center of his theory of psychology. Hobbes was careful not to claim that all of thought was verbal. Non-linguistic creatures could think, but Hobbes insisted that those mental processes where generality and orderly concatenation of thought are involved require the use of internal linguistic means. This remains to this day a highly controversial claim, although we will see later on in this book that many cognitive scientists today tend to agree that the emergence of language in the species played a key role in reshaping our cognitive landscape and altered our modes of thought in a dramatic fashion (see especially chapter 8).

What is clearly the case is that some of the core properties of language define those features that led Marcus⁸⁰ to talk of our mind as “the algebraic mind” and that have proven extremely difficult, not to say impossible to implement in the sort of generic neural network simulations that were popular in the 1980s. These are:⁸¹

- 1 the “type–token” distinction, or the distinction between a kind and an individual: between dogs in general and a particular dog (e.g., Fido);
- 2 compositionality: the fact that the meaning of the whole depends on the meanings of the parts *and the way they are combined* (the well-known distinction between the boring news “dog bit man,” and the much more newsworthy “man bit dog”);
- 3 quantification (or the binding of variables): the ability to provide the right set of brackets for elements like “for all x , . . .” and thus distinguish between⁸² *everyone in this room speaks at least two languages* and *at least two languages are spoken by everyone in this room*; or between *beavers build dams* and *dams are built by beavers*;⁸³
- 4 recursion: the capacity to embed one thought inside another: *John thought that Mary believed that Peter might have said that Chris was thinking that . . .*

Hobbes rightly saw these features of human cognition as essential. By seeing them as modes of linguistic computation, Hobbes was among the first to recognize the centrality of linguistic processes to our modes of thought.⁸⁴

To sum up this section, I hope to have convinced the reader that Cartesian linguistics offers the outlines of a productive research program that touches on the nature of linguistic knowledge, its acquisition, and its use – the first three questions of the Chomskyan linguistic program that find interesting parallels in the ethology literature (though, to be sure, the flexible character of linguistic use makes it quite unlike the capacities studied by the ethologists). It may be worth pointing out that the Cartesians did not ignore the question of neural implementation (the mind–body problem) or the question of origin. The rationalist-romantic literature is replete with proposals (some of them quite modern sounding) concerning the origin of human language, for the Cartesians knew that questions of origin could illuminate the nature of the cognitive faculty in question.⁸⁵ Like us, the Cartesians also knew that somehow the mental is the product of the brain,⁸⁶ but how the mental and the neural could be unified was to them as much of a

mystery as it is for us. Here too, the Cartesians have an important lesson to tell us: though the mental cannot be reduced to the physical, the mental should not be banished from rational/scientific discourse, as mental structures are key to rendering human cognition intelligible.

It is clear that the Cartesians had the right intuitions and the right set of questions; the decline of cognitivism during the behaviorist era is one of the most unfortunate developments in the history of science, and it is fitting that more and more studies go back to the lessons of the natural philosophers in search for fresh perspectives on a variety of issues. As Michael Wheeler wrote,⁸⁷ never underestimate Descartes.

5 Mathematical Foundations: Turing and Shannon

As we saw in the previous section, the idea that thinking is computing has a long history. Hobbes emphatically argued that reasoning ought to be captured by machine-like processes (“reasoning,” for him, was “nothing but reckoning”),⁸⁸ and, inspired by him (and others, such as the great Catalan philosopher Ramon Llull), Leibniz tried to develop a *calculus ratiocinator*, a calculating machine that can be seen as the grandfather of our modern pocket calculators. Later on, the British mathematician Babbage designed an engine for calculating accurate mathematical tables,⁸⁹ but it was not until the twentieth century that we gained a deeper understanding of what it means to say that something is a computer. Thanks to Alan Turing, Emil Post, Kurt Gödel, Alonzo Church, and other giants of logic and mathematics, computation became an object of mathematical thought, which in turn made it possible to achieve a scientific understanding of mental computation.⁹⁰

5.1 Turing

I will confine my overview to the work of Turing, which provided a formalism that expressed all the essential properties of what a computer does (in those days, the word “computer” referred to a person who engaged in the act of computing), of what could be computed in principle. This abstract mechanism is known as the Turing machine.⁹¹

A Turing machine is an extremely basic abstract symbol-manipulating device meant to capture the informal notion of “effective method” in logic and mathematics, and to provide a precise characterization of an algorithm (“mechanical procedure”). (It is important to stress that a Turing machine is a bit like a thought-experiment, an abstract model capturing the essential, logical structures of any given computation; the machine is a mathematical object; no actual, physical machine was meant to be built along these lines, even if it ended up inspiring engineering projects such as the modern computer.)

Turing was led to think about the notion of computation and computability in the context of a challenge formulated by the great mathematician David Hilbert in 1900.⁹² Hilbert asked whether it would be possible to devise a process according to which it can be determined in a finite number of operations whether a purely logical assertion (say, a mathematical statement) is true or false. Turing took on the challenge (and for the record, showed that no such process could be devised for mathematics), and in so doing gave us a precise characterization of what it is to compute.

The concept of a Turing machine developed by Turing is based on the idea of a person (a “computer”) executing a well-defined procedure to achieve a certain result. Here was, in his own words, Turing’s intuition:⁹³

Computing is normally done by writing certain symbols on paper. We may suppose that this paper is divided into squares like a child’s arithmetic book . . . I assume that the computation is carried out on one-dimensional paper, i.e., on a tape divided into squares. I shall also suppose that the number of symbols which may be printed is finite . . .

The behavior of the computer at any moment is determined by the symbols which he is observing, and his “state of mind” at that moment. We may suppose that there is a bound B to the number of symbols or squares which the computer can observe at one moment. If he wishes to observe more, he must use successive observations. We will suppose that the number of states of mind which need be taken into account is finite . . .

Let us imagine that the operations performed by the computer are split up into “simple operations” which are so elementary that it is not easy to imagine them further divided.

A Turing machine thus consists of:

- 1 a “tape”, which is divided into cells, one next to the other. Each cell contains a symbol from some finite alphabet. The alphabet consists of symbols (e.g., a special blank symbol (B), 0, and 1). The Turing machine is always supplied with as much tape as it needs to perform its computation. A tape thus looks like this:

...	0	B	1	1	1	B	0	0	B	1	...
-----	---	---	---	---	---	---	---	---	---	---	-----

- 2 a “head” that can read and write symbols on the tape and move the tape left and right one cell at a time;
- 3 a table of instructions that, given the state the machine is currently in, and the symbol it is reading on the tape, instruct the machine to do the following: (a) either erase or write a symbol and then (b) move the head one step left or one step right, and then (c) assume the same state or a new state as prescribed;

- 4 a state register that stores the state of the Turing table. The number of different states is always finite and there is one special start state. Turing defined this as a note of instructions to preserve the computation of the computer who is working in a “desultory manner” (the counterpart of the state of mind).

Note that every part of the machine (its state, alphabet, and set of actions) is finite and discrete; it is the potentially unlimited amount of tape that gives it an unbounded amount of storage space (memory). (Needless to say, this is not meant to be a realistic model of exactly how the brain functions. For one thing, the brain has no infinite tape. The point is, to the extent the brain computes, one must be able to construct some machine that can mirror/model what the brain does, and that such a machine will fall within the general description just given).

Consider a specific example of what a Turing machine could do.⁹⁴ To create the set of positive integers, we would have to start with a tape that contains only a symbol for 1 (say, a vertical stroke |); let's agree that the symbol for two is ||, for 3, |||. At the beginning, the head of our toy Turing machine is at a cell containing the symbol |. All the other cells contain the symbol 0.

The set of instructions we provide to the machine will be as follows:

- 1 When in S_1 and reading '|', move to the right and stay in S_1
- 2 When in S_1 and reading '0', write | and enter S_2
- 3 When in S_2 and reading '|', move to the left and stay in S_2
- 4 When in S_2 and reading '0', move to the right and enter S_3

Note that the table of instructions does not specify any rule for what to do in S_3 ; this means that the machine will halt (stop) if it reaches that state. Its computation will terminate. Note also that some rules specify a transition step to the state the machine is already in (“when in S_1 and . . . , stay in S_1 ”), that is to say, a machine does not have to change its state at every step.

Here is what our toy machine will do, step by step. At the beginning, rule (1) applies: the head of the machine reads |, and it is in the initial state S_1 , so it stays in that state and moves to the right. Now it is in S_1 , but it reads 0, so it writes | and enters S_2 (rule (2)). Now it is in S_2 , and reads | (the | it printed in the preceding step), so rule (3) applies. Therefore, it stays in S_2 and moves left. Now it is still in S_2 and is still reading |, so it stays in S_2 and moves left again. Having done this, the machine is still in S_2 , but it is reading 0, so rule (4) applies. Therefore it moves right and transitions to S_3 . Now the machine stops. It's back where it started but now the tape contains ||, which is the symbol for 2. If we run the machine again, it will end up at the cell it started but with ||| written on the tape, which is the symbol for 3, and so on.

This simple machine carries out the successor function; each time it is run it gives our symbol for the number that is the successor (next number) of the number whose symbol was on the tape before we started the machine. Note also that this is a maximally simple instance of a “recursive” procedure, a procedure that

can take its outputs as inputs, a loop-like device that can run up to infinity – an important notion in the context of human language, which as Humboldt put it enables man to “make infinite use of finite means.”⁹⁵

In so far as we know (this is known as the “Church-Turing thesis”),⁹⁶ the specifications given by Turing for his machine are all that is needed to perform any kind of computation (that is, whatever the brain does when it computes, the result can be captured by some Turing machine). No addition to the machine is needed to achieve more powerful results (though such additions may speed up the computation). Limiting the amount of tape, or reducing the number of possible actions, of course, limits the scope of what the machine can compute. We will see in Chapter 4 that Chomsky⁹⁷ put this to good use to show how different machines, with a varying amount of memory, symbols, instructions, etc., could model/compute the range of dependencies that are available in natural languages.

After Turing, whether or not any given mathematical function can be computed amounts to whether a Turing machine, set going with data and an appropriate program of computation (set of instructions) will ever get stuck (the technical term is “halt”) while trying to compute this function (a machine halts when there is not exactly one instruction rule specified).

At this point it may be useful to stress the limits of computation, and machine-like behavior.⁹⁸ Descartes is often portrayed as the archenemy of the human mind as machine – not because he was opposed to the mechanization of mind (quite the opposite, since he took mechanization to be the standard of intelligibility), but because he saw that some of man’s psychological capacities went beyond what could reasonably be achieved by a machine. In particular, he identified reason, common sense, and the sort of creativity displayed in language use as beyond the reach of a “mere” machine. The reason for this is very clear: for Descartes, machines were collections of special-purpose mechanisms (today, we would call these modules (more on which in Chapter 8)), and no single machine could incorporate the enormous number of special-purpose mechanisms that would be required to reproduce unbounded, creative, human-like behavior, or respond appropriately to the input while still being “detached” from that input (“stimulus-free”). These were the limits of the mechanization of the world’s picture, and they have remained insuperable problems for modern cognitive scientists. Three centuries later, Turing pointed out that mental intuition seemed beyond the reach of the computable (that is, beyond the reach of any Turing machine).⁹⁹ More recently, Fodor¹⁰⁰ expressed a similar skepticism, when he reviewed the achievements of modern cognitive science, pointing out that some aspects of cognition (such as inference to the best course of action, our guesses, and gut-feelings) remain outside the scope of modular (machine-like) characterization. So, to the extent we can talk about the mechanization of the mind, it is important to bear in mind that it remains a partial success (much like the mechanization of the world, as Newton showed, much to his dismay).¹⁰¹ But this should not be cause for despair: the fact that we can point to some limits in our understanding means that we have made progress, but it shows that one must be wary of grandiose statements pertaining to how the mind works.

5.2 Shannon

In standard histories of modern cognitive science, Claude Shannon's work¹⁰² on the mathematization of the notion of information occupies as central a place as Turing's work on computation. For the perspective that I want to convey to the reader of this book, this is not the case. One of the early results of Chomsky's work in linguistics was in fact to show how limited models of language based on information theory were (more on this in Chapter 4),¹⁰³ but I want to introduce Shannon's precise characterization of the notion of information because, as Gallistel has shown in a number of works,¹⁰⁴ such a characterization can illuminate the nature of learning, and show how wrong the behaviorists were when it came to what makes learning possible.

Like Turing did for computation, Shannon provided a rigorous characterization of the notion of information. Shannon's definition of information is, I believe, quite intuitive. For him, the amount of information conveyed by a signal is measured by the amount by which the signal reduces the receiver's uncertainty about the state of the world. If you already know that the first letter of my first name is "c", my telling you that the first letter of my first name is "c" tells you nothing. It conveys no information. But if you didn't know this before, then, of course, my message would be informative. This simple idea has non-intuitive mathematical consequences, though. Specifically, it implies that signaling (conveying information) presupposes prior knowledge on the part of the receiver. Recall that information is measured in terms of the difference between what you already know and what you don't. This means that for you to learn something, you must have some prior knowledge against which to measure the thing learned. That is, if the mind of the newborn baby were genuinely a blank slate, with no expectations about the sort of thing she could encounter, then any signal would convey an infinite amount of information. But according to Shannon's theory of information, no signal can convey an infinite amount of information in a finite amount of time, so no information would be conveyed, and no learning could take place. Put differently, in order for us to acquire information about the world from our experience of it, we must have built into our signal-processing capacity some structures/symbols ("representations") of the range of possibilities that could be encountered in the world. You may not know in advance what the first letter of my first name is, but you must know that it must be a letter drawn from a finite alphabet. That's already something (you know it's 1 of 26 options). Likewise, a thermometer will only be able to register information about external temperature if it is built for this specific purpose, if it has some notion of temperature built into it.¹⁰⁵ It is because we have some linguistic structure built into us that we can experience language; likewise, it is because we have no notion of infrared color built into us that infrared color remains inaccessible to us. What we can learn depends on what we are predisposed to learn. Our environment is specified by our biology. It is the organization of our brain that sets the limit on the information that we could in principle get from the signal, that is, what we could possibly learn.

In his mathematical characterization of information transmission, Shannon coded the fact that for communication to work efficiently, the meaning of the message was not the essential aspect (Shannon deemed meaning quite irrelevant for his purposes, in fact), but rather the fact that the message had to be selected from a set of possible messages. For this to work, the receiver had to know what the set of possible messages was, even prior to being able to adjust the relative likelihood of the different messages. If I don't tell you in advance that you must run only if I come out of the building wearing a red shirt, you won't know what detail to attend to when you see me exiting the building; you will be at a loss, unable to act appropriately.

The following example, taken from the linguistic domain, may help bring out the importance of this key point.¹⁰⁶ Consider the Chinese youth trying to learn the fact that the word *tian*, meaning "day", means "every day" when repeated (the technical term for this is "reduplication"): *tiantian*. There are no doubt various factors that the child must pay attention to when her caretaker produces the crucial utterance, but, more importantly in the present context, there are lots of details that she has to ignore: she must ignore the specific time of day at which the utterance was produced, the location of the caretaker, the clothes she is wearing, etc. – in short, she must pay attention to strictly linguistic factors. There are even linguistic factors that she has to ignore, such as the fact that the word being repeated begins with the sound *t*, or the fact that the word is monosyllabic, or bears a certain tone, or the fact that it refers to a time-related concept. All of that potential information must be ignored, lest the wrong reduplication rule be entertained (and we know from looking at the acquisition of reduplicative patterns that these potential errors are never made).

Chomsky¹⁰⁷ is fond of mentioning the following passage from C. S. Peirce, where the basic point at issue is made clear:

Besides, you cannot seriously think that every little chicken that is hatched, has to rummage through all possible theories until it lights upon the good idea of picking something up and eating it. On the contrary, you think the chicken has an innate idea of doing this; that is to say, that it can think of this, but has no faculty of thinking of anything else. The chicken you say pecks by instinct. But if you are going to think every poor chicken is endowed with an innate tendency toward a positive truth, why should you think that to man alone this gift is denied? I am sure that you must be brought to acknowledge that man's mind has a natural adaptation to imagining correct theories of some kinds.

In forming the idea of picking up food, and not countless other ideas it could hit upon, Peirce's chicken resembles a child learning the rule of reduplication in Chinese. Specifically, what they share is that the number of hypotheses they consider is restricted. It's more than just hitting upon the right thing, it's also moving away from the wrong sort of generalizations.

As Andrew Nevins has pointed out,¹⁰⁸ when we attempt to simulate the learning of a specific task, by, say, building a neural network that would simulate what the

brain does (as many cognitive neuroscientists do), it is important to remember that what we built into the network to give it a head start is just half of what's necessary if learning is to be successful. Equally important is what is not built into it. What allows neural networks to succeed, when they do, is not what they have been built to bring to the task, but rather what they are specifically built *not* to bring to the task. If neural networks kept track of every property inherent in the data, they would never be able to make any generalizations, let alone the right generalizations.

Nevins goes on to note that the function of what is called Universal Grammar/ the language faculty, then, is not really to provide a grammar, but rather to provide a set of constraints on what can and can't be a possible grammar. It's what we might call a selective (as opposed to a constructive) grammar.¹⁰⁹ It's the main difference between a traditional grammar book (which specifically tells you what's possible), and the sort of grammar that cognitive scientists posit, which tells you what's not possible and lets you explore the options that remain available. Nevins¹¹⁰ likens this state of affairs to what computer scientists call a "British Museum search," because it is likened to trying to find a pottery shard in the British Museum by simply looking through the whole museum. Though hard, the task is made easier by the "British Museum" part of the "British Museum search" (say, because it excludes other museums; that is, it leaves a lot open, but it excludes irrelevant domains of search, and automatically dismisses irrelevant sources of generalization.)

To sum up this section, as Randy Gallistel has correctly pointed out on numerous occasions, if we take the brain to be an organ of computation, as modern cognitive science does, and if it computes information it gets from the environment (meaning: it learns), then we are already in nativist territory, for there cannot be any information processing without any priors. Put another way, the brain could not run without a non-trivial amount of genetically specified structure about the world. Learning would be a total mystery if the brain were a blank slate; the external world would be lethally chaotic. Once this point is recognized, the essential questions of cognitive science are: what representations does the brain compute? From what data (external stimuli) does it compute them? How does it compute them (what are the specific computations it performs)? And how does a given representation get translated into observable behavior?

6 A Brief Summary

Under Chomsky's impetus, the task of figuring out exactly what sort of mental structures are deployed in the course of language use has always been high on linguists' research agenda. By paying attention to the range of computational manipulations that can be performed over these mental structures, we have learned quite a lot about language and how it can be learned.

David Marr¹¹¹ advocated a similar approach in the domain of vision, stressing that questions of neural implementation must be postponed until one is clear about what it is that the brain must compute. Other domains, such as music cognition or our faculty of moral judgment (our “moral sense”), have begun to follow the same path (see this book’s Epilogue), and no doubt research on other areas of cognition will follow suit.

The advances in mathematics in the area of computation and information have made it abundantly clear that rich inborn mental structures, highly specific to the tasks at hand, and the mechanisms/rules giving rise to them will have to be posited if one is to understand what the brain does – as the Cartesians were the first to stress, and as the ethologists discovered when studying lower animals.

This nativist stance advocated by Chomsky grew out of his recognition that learning was a computational problem. In Chomsky’s view, the learner must compute from data a highly specific mental structure that can then be put to (mental) use, the same way the ethologists thought of the animal as equipped with specific data processors. Now that the research agenda was clear, and equipped with the rigorous notions provided by mathematicians, cognitive scientists were ready to make some of the most essential aspects of our cognitive feats intelligible.

3

How the Mind Grows: From Meno to Noam

1 Gavagai!

Imagine you are taken to one of the most remote corners of the world, not knowing where exactly, not knowing who lives there, and what language (if any) they speak. After several hours of aimless wandering you meet one of the natives. He looks you up and down, silently, for what feels like an eternity. All of a sudden, his facial expression changes, and he fixes his gaze on something right next to your left foot, uttering something that sounds like “Gavagai!” You follow his gaze, and see a rabbit-like creature scurrying around.

Try to guess what Gavagai means. If you are like me, you will entertain a few possibilities, like “Watch out!”, or “Don’t move”, or “Run!”, or “Rabbit”. But how did you come to entertain these few possibilities and not a million others that the native may have in mind, like “OK, let’s go now!”, or “Incredible!” or – if you thought of the “rabbit” option, why not “undetached rabbit-parts”?

This little scenario is a variant of one told by philosopher Willard v. O. Quine¹ to stress that for any hypothesis you may entertain there are infinitely many that could be entertained and, for many of them, no amount of evidence would be able to make you choose among them. Think of what it would take to tease apart “rabbit” and “undetached rabbit-parts”; whatever is true of rabbit will be true of undetached rabbit-parts. Nothing out there will tell you to reject the undetached-rabbit-parts option. But it takes a philosopher to think of undetached rabbit parts. Yet, if you think about it, the Gavagai situation is one we all faced as children trying to acquire the meaning of words. How did we manage to guess that elephant refers to that big gray animal with a long trunk? Because someone pointed at the animal and said “elephant” (this was Locke’s suggestion²)? But how did you know what exactly was being pointed at? Surely the finger couldn’t point at the whole elephant, it was your cognitive bias that interpreted the act of pointing in that way. And how about words for less concrete/visible objects: words like *think*, *nothing*, *wait*, *digest*, etc.?

The list is endless, and yet, any normal child acquires all these words, and many, many others, sometimes (early in life) at the rate of several an hour. What's more, if you are like me, you don't remember suffering too much during this massive acquisition process. You probably didn't think twice about it, the process felt automatic. It just happened. You couldn't help it, much like (barring pathology) you can't help seeing things when you open your eyes (even things that are not "there," as in the case of visual illusions). You automatically interpret the environment in specific ways, you mentally structure it, you impose your cognitive mold onto it.

The Gavagai story provides yet another illustration of the fact that we have certain cognitive biases. I already provided a few examples of this sort of bias in Chapter 1, and in Chapter 2 I told you that such biases must be expected; all animals are biased in some way,³ that's what biology does for you. That's not special; what's special is the sort of biases that distinguish us from other species. Language largely contributes to the specific way we parse the environment, the way we "see" or "think" the world. In this chapter I would like to examine these cognitive biases a bit more closely, reinforcing and amplifying some of the ideas from the first two chapters of this book. I think it's important to provide as many examples from as many domains as possible to establish the existence of cognitive biases because our common sense all too often leads us to ignore them. I also want to introduce a few considerations bearing on the way such biases emerge in the child. That is, I want to touch on issues of psychological development and maturation, and show how the human mind grows from an initial state to an adult state. Once again, language will be the domain from which I will draw my examples, but the general conclusions I will reach go well beyond the linguistic domain. They in fact pervade the cognitive world.

2 The (Radical) Poverty of the Stimulus

By the time a child goes to kindergarten, virtually all of the key steps of cognitive developments, including language acquisition, have already taken place. Children accomplish these things all on their own, without specific instruction. Plato was perhaps the first to grasp the significance of this fact. In a celebrated dialogue, the *Meno*,⁴ he used Socrates to show how we intuitively know things we were never told before. Plato made this point by letting Socrates interact with a slave boy (i.e., someone to whom instruction was denied). By asking a few well-chosen questions, Socrates got the boy to express mathematical truths. It's as if the boy knew these truths all along, for the answers came so naturally to him. I have already used the same strategy with you, the reader, in Chapter 1. I didn't use mathematical concepts, but instead linguistic concepts. I was able to show that you know that *flying planes can be dangerous* can have two interpretations, but *flying planes is dangerous* doesn't. No special instruction was needed. You may not know what a wug is, but if I tell you that there is more than one, then you know that there are . . . "wugs."⁵ Likewise you know that if I say that *everyone didn't show up at the party*, I could

mean that “not everyone showed up at the party” (the other interpretation, that no one showed up, is actually hard to get). I could multiply these bits of unconscious linguistic knowledge at will. Like Socrates, if I ask you the right questions, you are bound to give me the answers that you may not have suspected were in you. But they must have been in you all along, for otherwise you would not be a speaker of English; you couldn’t have become a speaker (a linguistic creature), period.

The gap between knowledge and experience really is considerably bigger than we tend to think. Plato’s specific claim that the things we know in the absence of specific instruction are the things we remember from previous lives is wrong, but the intuition behind it is surely right. How else could we cope with what Chomsky has called the “poverty of stimulus,”⁶ the fact that what evidence the environment provides us with is too fragmentary and ambiguous to be able to account for the sort of knowledge we spontaneously manifest?

Many cognitive scientists to this day take the poverty of stimulus to be a hypothesis in need of strong arguments in its favor, but, to me, the poverty of stimulus is a fact about the world (one could also speak of a poverty of stimulus in the case of growing arms, digesting, reacting to viruses, etc.). One of the most striking things we have learned over the past 50 years is how severe the poverty of stimulus can be and how robust cognitive development turns out to be. I am here referring to instances of what is sometimes called language creation, or language acquisition in the near-absence of experience.

Perhaps the most famous case of language creation comes from the emergence of creole languages.⁷ To understand what creoles are, we must first understand what pidgins are. Pidgins are rudimentary systems of communication that emerge in particular types of language-contact situations, often arising from sad circumstances like population displacement, plantation economies, slave trade, interethnic exchanges, etc. Notorious among such situations are the plantation economies in the Caribbean during the seventeenth to nineteenth centuries, where African slaves furnished most of the labor force for Europeans. Linguists call the languages of the socially less-privileged groups in these situations the “substrate” languages, and the language of the more-privileged group is called the “superstrate” language. In such circumstances, a new linguistic variety, a *lingua franca*, invariably develops to bridge the communication gaps between speakers of mutually unintelligible substrate and superstrate languages. Such new varieties, developed and used by adults, are called pidgins. Pidgins are, by all accounts, very elementary, reduced, simplified systems of communication, with limited vocabularies.

Of supreme interest is the fact that the offspring of pidgin users go way beyond the properties of pidgins to which they are exposed; what they develop is a full-fledged language, a much more stable linguistic system that is called a creole, which very rapidly becomes unintelligible to the older generations. Whereas pidgins clearly fail to display the structural intricacies of natural languages, creoles behave like regular languages. Furthermore, quite a few structural characteristics are shared across creoles even if they emerged from distinct pidgins. Obvious lexical differences aside, a creole that grew out of the contact of, say, French and Gungbe displays interesting

similarities with a creole that emerged out of the contact of, say, Dutch and Ijo. For example, all creoles have a basic Subject Verb Object word order, even if some of the languages they grew out of have a basic Subject Object Verb word order. Likewise, no creole allows subjects to be left unexpressed (“left” for “he left”), whereas some of the languages they grew out of do (Spanish, Portuguese). And, all creoles have preverbal particles (what you may have been taught to call “auxiliaries” or “helping verbs,” such as *will*, *could*, or *must*) to express tense (past/present/future), mood (possibility/obligatoriness), and aspect (completed action/ongoing activity).

The most plausible interpretation of this fact, advocated by Derek Bickerton in his seminal work,⁸ is that this is our cognitive (linguistic) biases at work. Children of pidgin users find themselves in a dramatically impoverished linguistic environment (the pidgin) and essentially let their cognitive biases guide them toward the development of a genuine language. The similarities emerge due to the way all human beings are shaped (“pre-wired”) cognitively. In other words, the language faculty acts as a mold that quite literally transforms the input and refashions it according to the stringent precepts of Universal Grammar. In some sense, creoles may wear some of the core properties of Universal Grammar (the initial biases) on their sleeves. Creoles also show that language development cannot be a matter of imitation, analogy, or rote learning. Rather, it must be viewed as a creative act, as Humboldt was fond of stressing.

Another equally dramatic instance of language development in the near-absence of external stimulus comes from sign languages. Descartes⁹ had already observed that deaf children born to hearing parents with no knowledge of any sign language spontaneously develop a system of sign (now known as “home sign”)¹⁰ when they communicate among themselves. Although home signs are somewhat rudimentary systems by the standard of normal languages, they still display many of the properties of natural languages, and they go beyond the range of expressions that other species can communicate. Interestingly, studies have shown that the spontaneous signs used by speaking parents that do not conform to the grammatical principles of natural language are systematically “ignored” by children (in the sense that they are never imitated).

In the context of sign languages, the emergence of two new sign languages – Nicaraguan Sign Language¹¹ and Al-Sayyid Bedouin Sign Language¹² – has received a fair amount of attention in the recent literature. Unlike spoken creoles which emerged before the cognitive revolution was underway, the two sign languages in question could be investigated as they emerged, by trained linguists/cognitive scientists. As such they provided unique laboratory experiments in which to study language emergence processes in real time.

In the case of Nicaraguan Sign Language, the founding of schools for the deaf in Nicaragua in the middle of the 1980s amounted to the creation of a community of deaf children, none of whom had previously been exposed to a standardized sign language (deafness was stigmatized in Nicaragua, as it still is, sadly, in many cultures). Remarkably, an indigenous sign language quickly emerged, created entirely by the children. Whereas older children showed varying degrees of fluency in the language, younger children increased the complexity of the language, and brought it to the

standard of a full-fledged language – suggesting that age of development makes a big difference, as is clear from our difficulty in mastering a second language late in life (I will return to this fact later on in the chapter).

As for Al-Sayyid Bedouin Sign Language, it came to worldwide attention even more recently, when a group of linguists published the first study of this sign language that emerged in the last 70 years, uninfluenced by any other language, and used by an entire deaf community (and a considerable proportion of the hearing population) in a village in the Israeli desert.

Cases like Nicaraguan Sign Language and Al-Sayyid Bedouin Sign Language, perhaps more than spoken creoles since there cannot be any issue of substrate and superstrate influence, demonstrate beyond doubt that there is something to language acquisition that is more than just statistical correlation of inputs. Such cases indicate that the central question for cognitive scientists should no longer be whether there is a predisposition to grow a language (a term I prefer to “language acquisition” because “acquisition” suggests an E-language perspective, since it gives the impression of grabbing something outside one’s brain), but rather, what this predisposition includes. For the scientist, such instances of language creation¹³ offer us experiments of nature (experiments that could not be intentionally performed for ethical reasons) where learners are forced to go beyond what they have evidence for in a way that is more spectacular than in the usual poverty of stimulus circumstances. Still, I want the reader to realize that, however dramatic, cases of language creation differ only in degree from what must be taking place every day, for every child, raised in normal circumstances.

To be sure, in some societies, parents talk a lot to their children, and speak more slowly, or with exaggerated prosody (characteristics of what is often called “motherese”); they also usually supply the names of particularly salient objects, actions, and concepts (though I again emphasize, what counts as “salient” depends on internal/biological factors). But there is no evidence that these practices provide crucial help in language acquisition. They certainly cannot be regarded as essential, since in some societies/cultures, infant-directed speech is known to be rare,¹⁴ with no harmful effect on language development.

Furthermore, it is also clear that even in societies where children receive a lot of linguistic input, the data contains little in terms of explicit useful instruction (what’s often called “negative evidence” – evidence that something is not possible). What’s more, even when the child gets corrected by an adult, all the evidence points to the child ignoring the adult (as any parent knows), as in this well-known passage:¹⁵

CHILD: Nobody don’t like me.

MOTHER: No, say “nobody likes me.”

CHILD: Nobody don’t like me.

...

(after eight repetitions of this dialogue:)

MOTHER: No, now listen carefully; say, “*nobody likes me.*”

CHILD: Oh! Nobody don’t likes me.

In some cases of normal language development, what is “acquired” is so rarely seen in the environment that it is a safe bet to assume that the phenomenon in question could not be robustly (that is, effortlessly, unfailingly, by every child) acquired. This is arguably true of the “subject-auxiliary inversion” rule in English questions used by Chomsky to illustrate the concept of structure-dependence.¹⁶ Every speaker of English knows without being told that the way you form a yes/no question (a question that can be answered by yes or no) is by taking a regular “declarative” sentence and switching the order of the subject and the auxiliary. This is trivially illustrated in the following example:

He can go → Can he go?

He is tall → Is he tall?

But this simple rule does not explain what happens in more complex cases, like forming a question from *The man who is tall is hungry*. Every native speaker knows that the right result is *Is the man who is tall hungry?* and not, *Is the man who tall is hungry?* But, like Socrates, we may ask how native speakers know this fact. (And they “know” it from the very start, for no child has ever been observed to mistakenly produce something like *Is the man who tall is hungry?*) It seems logically plausible to assume that when there are several auxiliaries in a sentence, the first one is moved in front of the subject to form a yes/no question. This is in fact what children may “learn” by listening to alternations like *John can be happy* → *Can John be happy?* (not: *Be John can happy?*) But this would give us the wrong option in the example at hand (*Is the man who tall is hungry?*). You may be tempted to say that children simply listen to alternations like *The man who is tall is hungry* → *Is the man who is tall hungry?* However, alternations of this sort are vanishingly rare in conversations. To the extent that examples of this complexity exist, they are probably too few for the child to form the right hypothesis on the basis of them. Somehow, like Socrates, we have to assume that the relevant bias to manipulate the right auxiliary comes from within.

Ever since Chomsky mentioned this straightforward example, cognitive scientists who try hard to resist Plato’s solution (irrational as this may seem to readers of Chapter 2) have devised ingenious ways in which the child may “learn” the right “subject-auxiliary inversion” rule from experience. Although I think all the attempts have failed so far, I would like to point out that even if one could find a way for the correct “subject-auxiliary rule” to be inferred from the input, it’s likely to prove too specific, for all the attempts I am familiar with are based on the assumption that the rule has something to do with the auxiliary. In a sense they are right, but in another sense, they are deeply wrong, for it’s not just the first auxiliary that cannot be manipulated in *The man who is tall is hungry*. Consider the fact that one can say: *It’s hungry that the man who is tall is*, but one cannot say *It’s tall that the man who is is hungry*. The correct generalization (that is, the correct cognitive bias) is thus more general: no reordering rule of the sort we have illustrated (yes/no questions, “it is . . .” statements) can involve elements that

are in what is often called a “dependent” or “modifying” sentence/clause, like *who is tall* in our example (a clause that can easily be omitted without leading to an unacceptable/incomplete sentence; compare: *the man is hungry* and *the man who is tall*). How this generalization could be gathered from the input remains to me a total mystery.

As a final example of the vastness of the gap between knowledge and experience, consider the result of investigations looking at the acquisition of “visual” words by blind children.¹⁷ Vision is an important source of contextual meaning during language learning. So it would be a safe bet to assume that the restrictions on blind children’s access to contextual information ought to pose problems. And yet, detailed studies of the acquisition path demonstrate that there is neither delay nor distortion in their language growth. For example, words describing visual experience, like “look” and “see,” are the first verbs to appear in their language – with meanings adjusted to the sense of blind “looking” with touch. For example, when a sighted 3-year-old is asked to “look up,” he will tilt his head upwards, even if he is blindfolded. A blind 3-year-old raises her hands instead. If told “You can touch that table, but don’t look at it,” the blind 3-year-old will only lightly touch the table. If you later tell her she can look at the table, she may explore all the surfaces of the table with her hands.

It’s highly unlikely that blind children are explicitly taught these meanings for these words, and one can be sure that they do not receive any visual input, so we are again left with Plato’s conclusion: they must have created what are very reasonable meanings for these verbs by themselves.

3 Lenneberg’s Program

Pointing to the lack of evidence in the input for a given cognitive rule is one way to convey the message that there must be something in our nature that accounts for some of the most elementary properties of our behavior. There is a complementary way to reach the same conclusion. It’s the strategy that Eric Lenneberg pursued in the context of language in a series of essays in the 1960s,¹⁸ culminating in his landmark publication, *Biological Foundations of Language*.¹⁹ Lenneberg’s strategy was this: If the biases we must posit to account for behavior are part of our nature, that is, if we are equipped with, say, a language organ, we expect this organ to grow like other organs of the body (think of the teeth we grow, think of walking, and so on). It should display a robust developmental schedule, it should require proper nutrition, it should (unfortunately) be subject to disorders, etc. Lenneberg set out to look for clues of this sort, and although many details were to be filled in after his tragic death, his basic assessment remained correct: the child’s path to language displays the hallmark of biological growth.

Pursuing a biological approach to our capacity to develop a language, Lenneberg said, means that we should expect:

- 1 certain anatomical correlates
- 2 a fixed developmental schedule, uniform across the species
- 3 the inability to suppress such a capacity
- 4 the impossibility of teaching it to other species
- 5 the existence of universal properties
- 6 some genetic correlates, and therefore, some deficits

Ideally, some of these conjectures ought to be tested experimentally (with deprivation experiments, genetic manipulation experiments, anatomical experiments, etc.), but for ethical reasons, these experiments can't be performed on humans. Accordingly, those cognitive properties that are unique to humans, like language, must be tested only indirectly, with the occasional accident of nature strengthening results reached by careful analyses and inferences. (In other cognitive domains, not specific to humans, such as visual cognition, experiments are performed on other species, and the conclusions extended to humans, if we have reason to believe that the cognitive systems are sufficiently similar. This is why we know much more about the biological roots of our basic visual system than we do about language, music, mathematics, morality, and so on.)

With these restrictions in mind, we can now turn to Lenneberg's findings. Perhaps his most significant finding is the existence of a critical (or sensitive) period associated with language learning. The notion of a critical period is well known from the ethology literature,²⁰ and it received a boost of publicity when Hubel and Wiesel won the Nobel Prize for work that established the relevance of the critical period for the visual system.²¹ The primary focus of Hubel and Wiesel's experiments concerned information processing in the visual system. In one experiment, done in 1959, they inserted a microelectrode into the primary visual cortex of an anesthetized cat. They then projected patterns of light and dark on a screen in front of the cat. They found that some neurons fired rapidly when presented with lines at one angle, while others responded best to another angle. They called these neurons "simple cells." Still other neurons, which they termed "complex cells," responded best to lines of a certain angle moving in one direction. These studies showed how the visual system builds an image from simple stimuli into more complex representations. These experiments also established the extreme sensitivity of specific neurons. Thus, it looks like the range of angles that an individual neuron can register is set by the cat's genetic program, though experience is needed to fix the precise orientation specificity.

Hubel and Wiesel went on to perform experiments to identify how individual neurons in an animal's visual system react to specific patterns in the visual field (including horizontal lines, vertical lines, moving spots, and sharp angles). They discovered that particular nerve cells are set within a few hours of birth to react only to certain visual stimuli, and, furthermore, that if a nerve cell is not stimulated within that time frame, it becomes inert – its window of opportunity (critical period) has closed. Thus, in several experiments, they showed that if a kitten spent its first days in a deprived optical environment (such as a tall cylinder painted only with

vertical stripes), only the neurons stimulated by that environment remained active; all other optical neurons (pre-programmed to pay attention to other features of the visual world) became inactive because the relevant neural connections (synapses) degenerated. As a result, the kitten never learned to see horizontal lines or moving spots in a normal way. If, however, the kitten was placed in the same deprived environment a few days after birth (after it had had the chance to receive enough input for horizontal lines, moving spots, etc.), deprivation was not harmful. Similarly, if the kitten was deprived of the relevant input right at birth, but released earlier than in the first set of experiments, the original deprivation didn't have the same effect, and the kitten could recover.

These experiments demonstrate three important things. First, learning appears to be a selective process: kittens (and humans) appear to be genetically equipped with a set of specifications that structure our experience at birth. In the absence of the relevant experience, the structures that are pre-programmed degenerate, and learning fails. Use it or lose it. This is quite important, and quite counterintuitive. More and more evidence points to the fact that learning is, in fact, forgetting.²² The things we learn are the result of brain connections that were kept active, and therefore strengthened, during development. The other connections, even though they were pre-established by our genetic program, atrophy if we don't make use of them. (This view of learning is a bit like the process of natural selection, and for this reason the term Neural Darwinism²³ is used in this connection.)

Second, each region of the central nervous system seems to be tuned to specific learning problems and the different (computational) solutions they require: serving vision, neurons in some areas of the brain are concerned with contours, directions, and depth. Serving speech, neurons in some other areas of the brain are concerned with sounds of different frequencies, etc. Given the now well-established high degree of specificity of neural sensitivity, it's hard to see how some can still hold onto the idea that the same neural apparatus could deal with all these very different learning problems.²⁴ The brain is no jack-of-all-trades, one-size-fits-all solution; it's more like a collection of experts that solve highly specific problems.²⁵

Third, in addition to being pre-programmed to pay attention to specific things in the environment, the expert neurons are given a tight schedule in which to act. Past the deadline, they are dead (or, at least, significantly weakened). This is what a critical period is all about (if a particular capacity does not disappear, but is significantly weakened, it is customary to talk of a "sensitive" period instead). Coupled with the fact that learning different things requires different computational experts/solutions, we expect that learning will take place in stages: first, one property will be learned; once that window of opportunity closes, another window opens up, then closes, and so on, until the organism reaches a mature stage.

In the context of language, one finds some of the most compelling evidence for a critical period. For instance, we have plenty of evidence that babies are sensitive to a wide array of sound distinctions, as early as we can test them – as you might expect if they are born equipped to cope with the variety of sound distinctions that the languages of the world make use of. Crucially, we also have evidence that the

ability to discriminate sounds that the specific language surrounding the baby does not make use of (but that other languages in the world use) dramatically decrease with age.²⁶ For example, Japanese, unlike English, does not make use of the sound distinction between /l/ and /r/ (adult native speakers of Japanese have trouble distinguishing between *belly* and *berry*), but Japanese babies show a clear sensitivity to this (and many other) sound distinctions (they begin losing it around one year of age).

Before looking at other pieces of evidence for a critical period in language learning, this may be a good place for me to mention that testing babies is challenging. The testing method that is most commonly used is the so-called habituation technique. It involves tracking eye movement, which assumes that babies turn their gaze when something new catches their attention, or keeping track of the sucking rate for the youngest babies, which assumes that babies suck on a pacifier more when they detect something new. The idea in both techniques is that the babies show a capacity to, say, distinguish between two sounds (such as /p/ and /b/)²⁷ on the basis of the fact that, after having been exposed to one sound for a while, they shift their attention when a new sound is heard.

Using ingenious techniques like these, it has been demonstrated that, at birth, babies can recognize their mother's voice and distinguish it from other voices.²⁸ The research was carried out by exposing the babies to two voices, their mother's and that of a stranger. One group heard the stranger's voice if they sucked hard, and the other, their mother's voice. Sucking less hard meant they heard the alternative voice. Babies chose to suck at the rate which meant they heard their mother's voice, demonstrating that they recognized it.

The same approach was used to see whether babies could remember a specific story that their mothers had read to them daily during the last three months of pregnancy. Again, in a similar way to the study above, the babies showed a memory for this story, compared with other stories.²⁹ This suggests that not only could babies remember their mother's voice, but they could distinguish a particular story read by her. They could also distinguish this story from another when both were read by a stranger. This suggests that the infants had learnt specific acoustic features of the story.

Research has also demonstrated a preference of newborn babies for the human voice over other sounds of similar pitch and intensity,³⁰ and a preference for sounds within the human voice range to sounds outside the human voice range.³¹ Other studies have shown that newborns can discriminate between some speech sounds, for example consonants such as /p/ and /b/ and vowel sounds.³² At two days they can also discriminate their own language from other languages, as well as a language with a similar prosody (intonational pattern) as their own from a language with a distinct intonational pattern.³³

Results of this sort³⁴ leave little doubt that not only human babies, but also human embryos, are biased for language.³⁵ While the results just mentioned indicate that human brains are fine-tuned to the sounds of language, it has also been observed that normal speech development necessarily goes through a series of language

production stages: first, a babbling stage, at about 6–9 months of age, followed by the one-word stage (9–14 months), then a stage where the child produces two-word strings (17–26 months), and eventually, a series of grammatical fine-tunings that I will touch on in Part II.³⁶

In the preceding paragraphs, I indicated that babies are attuned to speech, but research on deaf communities has established that the same fine-tuning and sensitivity exists in the realm of sign. Signing babies go through the same stages, including a hand-babbling stage.³⁷ Cognitively speaking, the difference between sound and sign appears negligible. It's the same language faculty at work.³⁸

Unfortunately, the powerful language faculty underlying the wonderful linguistic developments in humans is, like all organs, quite fragile. It depends on the right sort of “nutrition” (read: environmental input). Deprived of that input, the language acquisition process can be disrupted in the same way the visual system was in Hubel and Wiesel's experiments. Several instances of such terrible deprivation have been reported in the literature.³⁹ The most famous such case is Genie.⁴⁰ Genie spent nearly all of the first 13 years of her life locked inside her room. She was the child of a mentally unbalanced father, who decided that she was profoundly retarded, and subjected her to severe confinement and ritual ill-treatment in an attempt to “protect” her. During the day, she was tied to a child's potty chair in diapers; at night, she was bound in a sleeping bag and placed in an enclosed crib with a cover made of metal screening. Her father beat her every time she vocalized, and he barked and growled at her like a dog in order to keep her quiet. He forbade his wife and son to speak to Genie. By the age of 13, when she was discovered, Genie was almost entirely mute, commanding a vocabulary of about 20 words and a few short phrases (nearly all negative), such as “stopit” and “nomore.”

At that point, she was subject to close attention by various psychologists. Of immediate relevance to this chapter is her inability to develop a full-fledged language. Despite intensive training, and remarkable progress in building up a vocabulary, as well as a desire to communicate, Genie never learned fully grammatical English and only went so far as phrases like “Applesauce buy store.” Although she was very eager to learn new words (and grew frustrated when the people around her failed to give her a special word for, say, a new color that she was attending to), she never developed a grammar.

Genie is sadly not the only feral child we know of. The ability of feral children to learn language on their return to human society is very varied. For most feral children from history, we don't have enough information to judge exactly how much language, if any, they might have been able to learn. But some, like Isabelle,⁴¹ have been reported to acquire normal language ability. The crucial factor seems to be when they were discovered and restored to an environment with the appropriate mental “nutrition.” If this happens before the very end of the crucial period (traditionally taken to be puberty), the damage can be undone. But, if, like Genie, discovery takes place after the onset of puberty (roughly, around the age of 12), the process is irreversible, as Hubel and Wiesel discovered when they were investigating the visual system.

In the case of Isabelle, she was discovered when she was 6. It is reported that her progress was dramatic: in two years she covered the stages of learning that usually take six years. (This is reminiscent of cases of delayed language production, which Albert Einstein is said to have suffered from. Sometimes children in normal situations, for unknown reasons, do not start talking until they are as old as 4; but when they do finally begin talking, their performance level is normal for their age.) Cases like Genie remind us that despite our innate bias to grow a language (and many other cognitive structures), one must never lose sight of the fact that normal language development requires direct experience with language to “trigger” the inborn ability to acquire grammar (the same is true of most abilities of animals studied by the ethologists).

The contrast between Genie and Isabelle demonstrates that language acquisition is clearly maturationally constrained, and that something like a critical period holds for language acquisition. Please note that when I say critical period for language, I am not being precise enough, for language is not a monolithic capacity – it is made up of various computational tasks associated with distinct critical periods, many of which close well before puberty is reached. But puberty can be taken to be the very end of this remarkable window of opportunity.

Aside from cases of feral children, support for associating the final critical period with puberty comes from the sharp contrast experienced by many of us trying to learn a second language “late” (for language-learning, even the teenage years are quite late!) in life. We all know that the task is nothing like acquiring a first language. The process is far more conscious, the difficulties numerous, and the end result, typically, disappointing. Many of us achieve some degree of fluency in a second language, but native-like fluency is rarely, if ever, achieved.

Perhaps the most dramatic case of second language learning I am familiar with is the case of the late MIT Professor of linguistics Ken Hale. Hale’s linguistic ability was the stuff of legend.⁴² He was able to master a new language in the course of a week, based on limited contact with a native speaker, and a good dictionary. He was known to lecture in many different languages, and was virtually flawless in each (so much so that many misidentified his country of origin). For Hale, it’s as if the critical period never ended. But for the vast majority of adults, no matter how motivated we may be, second language acquisition is a source of frustration. There is nothing like a first language. No matter how often one is forced to use a second language, certain cognitive tasks appear to demand the use of a first language (try doing arithmetic in a second language, even one you are fluent in; you’ll see how once you get beyond simple numbers and computations, you will automatically switch to your first language).

Let me end this brief discussion of second language learning by pointing out that learning a second (or third, or n th) language before reaching puberty is quite different. A large proportion of the world’s population is multilingual (come to think of it, if we resist the traditional distinction between a language and a dialect, I suspect all of us speak at least one dialect and one standard language, so all of us are to some extent multilingual), and we know of no upper limit (apart

from social/political/practical ones – after all, there must be enough input for each language for the child to learn them all) to the number of languages an individual can acquire during the critical period. Children seem remarkably adept at keeping track of them all, and keeping them distinct.

All in all, the evidence for a fixed developmental pattern and for a (series of) critical period(s) seems very strong, and each new piece of evidence reinforces Lenneberg's original argument for a biological foundation for language acquisition. Lenneberg⁴³ was also right when he suspected that for cognitive capacities like language, one would find anatomical correlates, computational universals, and (at some level) virtually no variation within a species, but a great gulf between species. All these characteristics are discussed in subsequent chapters.

4 Wrapping Up

This chapter provided some of the very best reasons we have to believe that language is a species-specific trait, with a strong biological endowment. The fact that all infants, barring pathology or unusual environments, acquire linguistic material at a fairly brisk pace, with little explicit teaching or corrective feedback, was already suggestive. Cases of language creation, where children appear to invent aspects of linguistic structure, as pidgins are transformed into creoles, or when deaf children streamline or embellish sign languages that have been awkwardly modeled by their non-natively signing parents, really indicate that humans come to the world equipped with a language-ready brain – what Lenneberg called “man's preparedness for speech[/sign].”⁴⁴

The experience of language used in a community is what awakens the language faculty. Once awakened, the faculty shapes and guides the growth of grammars in children – hence the regularity, speed, and ease of language acquisition; hence the critical period effects, hence the commonalities of structures across languages, and hence (as we will see in Part II) the limits on language variation.

One consequence of the strong role played by *internal* factors in the case of language is a course of development that appears relatively uneventful to outside observers, who may think that language, given its size and complexity, should require more effort than is actually witnessed (as it does when we try to learn a second language as adults). Indeed, many people who are exposed to this biological perspective on language reject it out of hand, as it doesn't leave much room for the role of culture that they intuitively ascribe to the nature and acquisition of language. But such a reaction should be avoided at all costs. After all, the examples of language creation and deprivation reviewed in this chapter show that what awakens language is the experience of language in a community. In this way language is different from the visual system, where the triggering experience is not dependent on social factors. But reliance on the community need not conflict with the label “biological.” As Jim McCloskey⁴⁵ points out (pushing an analogy that Lenneberg⁴⁶ used as well):

nobody would deny that learning to walk is a biological process, the result of our evolutionary history, and the outcome of complex physio-chemical events. It is clearly under the control of a genetically-determined bio-program [hence the fixed developmental schedule, critical period, etc. associated with it]. And yet walking too can be said to be a social and cultural practice, with important symbolic and cultural functions. We learn to walk under the eyes of our parents and siblings, and their encouragement and guidance – though it does very little biologically – seems to us and to them to be central in our learning to walk. Furthermore, as soon as we have learned to walk, we develop particular ways in which we put this capacity into practice, ways that become culturally significant and distinctive [and convey information]. There is no conflict between the social and the biological, and the same should be true for language.

One must bear in mind that nature and nurture are false dichotomies. Genes in isolation have been said⁴⁷ to be among the most impotent and useless materials imaginable. To function, they require signals, some of which originate outside the organism.⁴⁸ As for the environment, the effective environment (what the ethologists call the *Umwelt*; cf. Chapter 2) is a function of the organism's perceptual orientation, which is affected in large part by genes.⁴⁹ What counts as a linguistic community for a human child surely looks like a different environment for a cat, or a pet fish. It's due to their different biologies, which in turn feed on the environment.

The difficulty with language is that we are so constantly aware of its social and cultural significance (everything around us reminds us of this fact) that its biological foundations are all too easy to miss. Although I think many linguists and cognitive scientists would agree that biology plays some role, many still subscribe to a view of human nature as plastic, ready to adjust to any environment, and endorse a view of learning that is as general as possible. But biologists are used to a different view. Since Darwin, they have recognized that the exquisite adaptation of organisms to their niches requires special-purpose mechanisms, fine-tuned to their tasks. It is this view that cognitive science must adopt, otherwise unique cognitive capacities like language are bound to remain shrouded in mystery.

PART II

Unweaving the Sentence

4

Mental Chemistry

1 The Hidden Texture of Language

As Nobel laureate physicist Frank Wilczek once noted,¹ science is based upon the evidence of our senses, and yet transcends them. For Wilczek this paradox is well illustrated by the theory of color. Science early discovered, through the experimental genius of Newton, that light “in itself” has a much richer structure than our sense of vision reveals. Physical color is not the same as sensory color; the former is more fundamental than the latter, even if the latter is what we actually see.

For me the paradox that Wilczek talks about is well illustrated by the theory of language that emerged in the context of the cognitive revolution.² Thanks to the experimental genius of Chomsky, we have come to understand that linguistic units (words, sentences, etc.) have a much richer structure than our perception of them reveals; or, I should say, our common sense leads us to believe, for it is our distinctive perceptual-cognitive apparatus that is responsible for us perceiving words and other abstract units in the speech signal, which in and of itself is a continuous signal. The structures we impose on this signal are the real stuff of language, what Andrea Moro has called³ its hidden texture.

The first thing to be clear about is the existence of discrete units in language: a sentence is not like a wave, it's made up of little quanta or atoms we call words (or, more technically, morphemes, the smallest bits of language that convey meaning; you may remember being drilled on these in school: trans-, un-, tri-, -ism, etc.). There are sentences of 4 words, 5 words, 6 words, etc., but no sentence of 4.5 words, or 6.3 words. Our mind structures the linguistic input in a digital form (as opposed to an analog form), and we call this property of language *discreteness*.

Next, linguists discovered that words in a sentence are not like beads on a string; the discrete atoms we call words combine to form molecules (known as “constituents” or “phrases”), whose arrangements are dictated by strict laws that ultimately give rise to specific meanings. If you are not yet convinced of the spontaneous structuring power of the mind, consider this passage from Lewis Carroll’s *Jabberwocky*:⁴

’Twas brillig, and the slithy toves
Did gyre and gimble in the wabe;
All mimsy were the borogoves,
And the mome raths outgrabe.

You don’t know what *brillig* or *toves* or *gimble* mean, but you can’t help process the lines I just quoted from as if *wabe* is a noun, and *gimble* is a verb, pretty much like the following (poetic effects aside):

’Twas dark, and the leafy branches
Did move and tremble in the wind;
All scared were the children,
And the big dogs barked.

The little words you recognized in the *Jabberwocky* passage, the *and* and the *the* and the *in*, helped you structure the unknown words like the little pac-man figures force you to see a triangle in the well-known visual effect (so-called Kanizsa triangle) in Figure 4.1.⁵ (With Lila Gleitman,⁶ I think this is exactly how children learn many of the words in their language, narrowing down their possible roles and meanings by using the little words around them. This is sometimes called “syntactic bootstrapping.”)




Figure 4.1 The Kanizsa triangle

The computational properties of language – its “syntax” in the broad sense of the term – can be studied at many levels (technically known as “levels of representation,”⁷ though perhaps a more transparent term would be “levels of structure”): at the sound level (phonology), the word level (morphology), the sentence level (syntax in the narrow sense), the level of meaning (semantics), and the level of use in specific situations (pragmatics). Taken together, all these properties and layers of structure constitute our knowledge of language.

2 Modeling Infinity

Perhaps the most interesting aspect of this knowledge is the property of infinity (underlying our creative aspect of language use). If there is one universal property in language, that would be it. Just like you know that you can always add one to any number and make it bigger, you know that you can always expand any sentence into a longer one (try “*x* said that . . .”, or “and . . .”).

Accordingly, a central task for linguists/cognitive scientists is to find a way to represent structure that allows for infinity.⁸ Due to the property of infinity we know that linguistic knowledge can’t boil down to a list (of behaviors/sentences). The brain is finite, so no infinite list can fit in it; on this point Chomsky is fond of quoting Humboldt, who noted that language “makes infinite use of finite means.” What must be found is an appropriate mechanism (a set of rules/instructions) that generates (in the mathematical sense of the term) an infinite array of expressions, just like the formula $f: \forall x \rightarrow x^2$ will take any number x and multiply it by itself. In other words, we must characterize our knowledge of language intensionally (by means of a general mechanism/rule), not extensionally (by merely listing all the instances conforming to the rule). (The intensional characterization is one of the meanings behind the “I-” in Chomsky’s “I-language” (and correspondingly, “I-linguistics”)⁹ coinage mentioned in Chapter 1.)

In his landmark *Syntactic Structures*,¹⁰ Chomsky discusses three successively more powerful kinds of mechanisms (and resulting structures) to model the range of expressions that our language faculty allows. Because it is such a seminal study I want to spend some time discussing it. We’ll see that understanding the limits of each model can teach us a lot about what we tacitly know about language.¹¹ Let me warn the reader that the presentation may look technical at times, but it is necessary. Like Turing did when he proposed what at first looks like a description of a computing machine (cf. Chapter 2), what we are trying to do is come up with a precise characterization of a specific algorithm/mechanism, and there is no substitute for formalism to capture the logical structure hiding behind familiar example sentences.¹²

2.1 Finite-state machines

The first model that Chomsky discusses¹³ is the simplest one that would allow for an infinite range of expressions. It’s called a finite-state automaton (or machine), and it was the sort of mechanism people looked at in the context of information science (cf. Chapter 2).

Finite-state machines are very simple devices consisting of (1) an initial state, (2) a finite number of additional states, (3) a specification of transition from one state to another, (4) a specification of a (finite number of) symbol(s) to be

“printed” when a particular transition occurs, and (5) a final state. Such machines can successfully model sentences like:

- (1) The pretty girl ran.


This could be represented via a finite-state machine as:



According to the machine, the sentence in (1) would have a very simple, linear structure, with words introduced one at a time, like beads on a string – pretty much like what you would do if you were to write down the sentence in (1) on a blank sheet of paper.


Such a machine can generate an infinite array of expressions if they are allowed to “loop.” For example, our machine can represent infinite variations of (3) via a device like (4).

- (3) The (very very very very . . .) pretty girl ran.



The machine in (4) captures what you would do if each time you write down *very*, you (as it were) went back and wrote the same word again and again until you decided to introduce some new word. The loop does the trick. But although finite-state machines effectively capture infinite expressions, Chomsky quickly noted that they run into problems in a number of circumstances. In particular, Chomsky showed that finite-state machines are formally incapable of modeling certain dependencies between expressions in a string. The main thing about finite-state machines is that they have no “memory” to work with: the only thing the machine “knows” at each step of the algorithm is what state it is in, and how it can get from that state to the next, printing one item at a time. Crucially, it is constructed in such a way that it does not remember the steps it took before; it lacks the ability to look at the tape in Turing’s machine, if you will. That’s what makes it run into trouble when it has to keep track of things, as I will now show you.

Consider, for example, the following artificial language (call it Martian). It’s a very simple language. It consists of two “words”: *a* and *b*. The finite-state machine discussed so far would have no problem capturing its “sentences”: *ab*, *aaaab*, *aabbbbbbb*, etc. All you need is this:



Now imagine a slightly different language, Martian2, that is similar to Martian (it has the same words, *a* and *b*), but different in that it requires its sentences to have the same number of *a*'s and *b*'s: *ab*, *aabb*, *aaabbb*, *aaaabbbb*, etc.

Finite-state machines cannot capture this. The best they can do is represented in (6):



The machine can generate an infinite number of *a*'s followed by an infinite number of *b*'s, but cannot make sure that there will be an equal number of *a*'s and *b*'s. It may turn out that the machine loops the same number of times for *a* and *b*, but this would be pure coincidence, and it wouldn't be Martian2, which requires the same number of *a*'s and *b*'s all the time. The reason for this limitation of finite-state machines is easy to identify. It's their lack of memory. They do not know what states they used to be in or how many times they have been in that state. They have no way to connect the "print-*b* state" with the "print-*a* state." Once they reach the "b-state," it's as if they're back to square one.

Of course, the next question that immediately arises in the face of this limitation of finite-state machines is whether this limitation has any bad effect on capturing real linguistic examples, and not just toy languages like Martian and Martian2. As it happens, examples that are impossible for finite-state machines to handle abound in natural languages. Think of cases where you say "either . . . or . . ."; for every *either*, there will have to be an *or*; of *if . . . then*; or every time you open a (mental) bracket, you'll have to close it: (. . . (. . . (. . . (. . .) . . .) . . .) . . .). Or think of your linguistic ability to let the subject trigger agreement on the verb (for instance, making it plural), even if the element that controls agreement on the verb can be arbitrarily far away from the verb: *The men [that John said that Mary believed that Sue saw] were* (not: was) *sleeping*.

One of the best examples of impossible dependencies for finite-state machines comes from my teacher Howard Lasnik,¹⁴ who got it from his teacher, Morris Halle, a key figure in the cognitive revolution. It's about the word *missile*. Countries that have *missiles* typically also have ways to counterattack other countries with similar weapons. They have *anti-missile missiles*. Setting aside technological limits, we can imagine countries having *anti-anti-missile-missile missiles*, *anti-anti-anti-missile-missile-missile missiles*, and so on. The pattern is quite clear, and is captured in the schema in (7).

- (7) $\text{anti}^n \text{missile}^{n+1}$ (n occurrences of the word *anti* followed by $n + 1$ occurrences of the word *missile*)

Finite-state machines fail to capture (7) as they are unable to keep track of the number of *anti*'s in order to make the number of *missile*'s greater by 1.

Every time there is some sort of dependency between an element x in a sentence and another element y that is not immediately before x , the finite state-machine will be at a loss. Words really aren't like beads on a string, coming one after another. Even if they allow for infinity, finite-state machines can't be taken to be adequate models for linguistic structures (though they can be adequate in some cases: as every child knows, when her mother says *Never, never do that again!*,¹⁵ which does not have the structure (never: (never do that again)) – that would be a positive command, not a threat (the two negations would cancel each other out).)

As the reader will have already realized, what we need is a more hierarchical representation for linguistic structure, which would allow us to keep track of a 's and b 's in *Martian2*, like this:

- (8) $\begin{array}{ccc} a & & b \\ & a & b \\ & & a\ b \end{array}$

This is the second model for linguistic structure that Chomsky examined in *Syntactic Structures*. It's called a rewrite-rule system (or a Context-free Phrase Structure Grammar). The name does not matter too much, but the properties of this machine are very important for language.

2.2 Rewrite rules

Like all machines, the rewrite-rule system starts somewhere (i.e., it has a designated initial symbol, typically: Σ), and a key device, the rewrite rule. The rewrite rule consists of one symbol to the left, followed by an arrow, followed by at least one symbol ($A \rightarrow B$) (the arrow stands for "rewrite as" or "consist of"). It also contains a procedure for moving from one step to the next known as a derivation, where exactly one symbol is being rewritten as another (sequence of) symbol(s) per step. The machine stops when there are no more instructions to rewrite a symbol as another.

As we did for finite-state machines, let us take one example and see how phrase structure grammars work by means of a toy example. Consider (9):

- (9) a. $\Sigma: S$
 b. $S \rightarrow ab$

What the machine in (9) says is start with the symbol S (9a), and rewrite it as ab (9b). There is no instruction to rewrite a or b as anything further, so the machine generates the expression ab .

To generate an infinite array of expressions of the sort required in Martian2, all we have to add to (9) is this:

(9) $b'. S \rightarrow aSb$.

The presence of “ S ” to the right of the arrow serves as an instruction to go back to the symbol to the left of the arrow and keep applying the same rewrite rule. (Ultimately, the rule in (9b) will have to be picked once to get out of the loop. Otherwise the machine would never stop cranking out aSb , $aaSbb$, $aaaSbbb$) Introducing the symbol to be rewritten into the output of the rewrite rule has the same effect as the loop in our finite-state machine. The difference between the two is that now we have a way to keep together the elements that form a dependency (the a ’s and the b ’s): they are both introduced by the same rule. Every time you apply a rewrite rule yielding an a , it is bound to yield a b (if you need two b ’s for each a , as in the *anti-missile missile* example, use the rule $S \rightarrow aSbb$).

The key property of a rewrite rule machine is thus the ability to pair up elements that can end up infinitely far apart (so-called unbounded discontinuous dependencies), like *men* and *were* in our example above (*the men that John said that Mary believed that Sue saw were sleeping*). Like the loop, the rewrite system is recursive (it applies to its own output), but it yields a different kind of structure. Rewrite-rule machines have the effect of creating hierarchical or embedded structures. They function like the Russian dolls that you are probably familiar with: you start with a big doll that, once you open it up, contains a similar-looking doll that, once you open it up, again contains a similar-looking doll. The similar-looking doll is the “ S ” in our toy machine. It ends when there is no more doll that opens up, no more S ’s to be rewritten.

This hierarchy-creating, embedding effect of rewrite rules is crucial for language, because in that cognitive domain hierarchy is of the essence. From now on, the reader should no longer think of examples like *Mary said that Peter thought that Sue believes that . . .* in terms of adding a sentence at the end of another sentence at the end of another sentence, but instead as adding a sentence inside another sentence inside another sentence, as in Figure 4.2. Only then will we be able to understand (i.e., offer a scientific account of) why *men* and *were* can “see





Figure 4.2 Embedded sentences

- (13)
- | | |
|----------------------------------|--------------------------------------|
| Step 1: S | |
| Step 2: NP VP | (via $S \rightarrow NP VP$) |
| Step 3: N VP | (via $NP \rightarrow N$) |
| Step 4: Mary VP | (via $N \rightarrow \text{Mary}$) |
| Step 5: Mary V S | (via $VP \rightarrow V S$) |
| Step 6: Mary thinks S | (via $V \rightarrow \text{thinks}$) |
| Step 7: Mary thinks NP VP | (via $S \rightarrow NP VP$) |
| Step 8: Mary thinks N VP | (via $NP \rightarrow N$) |
| Step 9: Mary thinks John VP | (via $N \rightarrow \text{John}$) |
| Step 10: Mary thinks John V | (via $VP \rightarrow V$) |
| Step 11: Mary thinks John laughs | (via $V \rightarrow \text{laughs}$) |

The workings of a rewrite rule system can be (and often are) represented via graphs known as tree diagrams. For example, (13) can be represented as (14):




Trees of this sort are ubiquitous in the linguistics literature. They are great visual aids, and I will make use of them later in this book.

2.3 Crossing dependencies


There is a lot more to say about hierarchy in language, but before I do so, I want to discuss the last big point Chomsky made in *Syntactic Structures*. It contains one of the truly greatest linguistic analyses ever.

Chomsky argued that although rewrite rules/phrase structure grammars are great at capturing unbounded discontinuous dependencies of the sort we discussed, some additional mechanism is needed to fully capture the notion of (mental) structure in language. Here, unlike in the case of finite-state machines, Chomsky¹⁶ didn't show that some relations could not be captured via rewrite rules in principle; rather, he argued that the necessary extension would render rewrite rules for language "extremely complex, *ad hoc*, and 'unrevealing'." Since the point of any formalism is to shed light on what looks like hopeless chaos, it stands to reason that we should be suspicious of a formalism that itself would be very complex.

The type of relation that rewrite rules/phrase structure grammars fail to capture neatly is technically known as a cross-serial discontinuous dependency. The discontinuous dependencies we have looked at so far all have a nested pattern (think of the way the structure would have been generated from the rewrite rules in (9): the outermost pair of *a* and *b* first, then the next pair popping up from the *S* between them, and so on):



Crossing dependencies look like:



The clearest and most famous example of cross-serial dependencies in language comes from Chomsky's celebrated discussion of the English auxiliary system in chapter 7 of *Syntactic Structures*. A detailed exposition of this argument would take me too far afield, so I will limit myself to offering a taste of Chomsky's reasoning. But I hope that the reader will appreciate the depth of the analysis.¹⁷

In addition to main verbs in sentences like *John sings*, *John laughed*, English makes use of auxiliary verbs that give rise to sentences like *John has sung*, *John will sing*, *John must run*, *John is running*, etc. There are basically three types of auxiliaries in English: the group linguists call modals (they express the "mood" of the sentence: *can*, *must*, *may*, *will*, etc.), and then two others, each in a class of their own, namely *have* and *be*. Multiple auxiliaries can be combined into sentences like *John will have left*, *John will be running*, *John has been running*, *John may have been running*, etc.

Chomsky's genius was to uncover significant generalizations when auxiliaries are combined, and express them in a compact, equation-like formula. The generalizations are the following (for each generalization I will provide an example sentence that violates it, which I will indicate by a * symbol in front of it, as is common in linguistics):

Generalization A: When a sentence contains a modal auxiliary (call it *M*), it is always the first verb-like thing after the subject in a declarative sentence. (**John has must be running*)

Generalization B: In the presence of *M*, no verbal element (auxiliary or main verb) bears agreement morphology that cross-references the subject (such as the -s on

has, signifying that it goes with a third-person singular subject). (**John must has been running*)

Generalization C: When some form of *have* and some form of *be* co-occur, *be* immediately follows *have*. (**John must be have running*)

Generalization D: The main verb is always the last element of the verbal sequence. (**John must have running be*)

Generalization E: If the main verb is the first verblike thing in the sentence, then it is inflected for tense (present/past), and for subject-agreement. It cannot appear bare, or in the progressive (-*ing* form) or perfect (participle) form. (**John run*)

Generalization E': If *have* is the first verblike thing in the sentence, then it is inflected for tense (present/past), and for subject-agreement. It cannot appear bare, or in the progressive (-*ing* form) or perfect (participle) form. (**John have run*)

Generalization E'': If *be* is the first verblike thing in the sentence, then it is inflected for tense (present/past), and for subject-agreement. It cannot appear bare, or in the progressive (-*ing* form) or perfect (participle) form. (**John be running*)

Generalization E''': If M(odal) is the first verblike thing in the sentence, then it is inflected for tense (present/past: e.g. *will/would*), but cannot appear bare, or in the progressive (-*ing* form) or perfect (participle) form. (**John musting have been run*)

Generalization E+: Whatever verblike thing is first in a sentence, it will be inflected for tense (present/past), and for subject-agreement. It cannot appear bare, or in the progressive (-*ing* form) or perfect (participle) form.

Generalization F: When a sentence has two verblike things and the first is some form of *have*, the second appears in the perfect (participle) form. (**John has being run*)

Generalization G: When a sentence has two verblike things and the first is some form of *be*, the second appears in the progressive (-*ing*) form. (**John is run*)

Generalization H: When a sentence has two verblike things and the first is an M, the second appears in the bare form. (**John must having been run*)

You may not be aware of all these generalizations about English verbal morphology, but all native speakers of English (unconsciously) rely on them when producing sentences.


Chomsky captured all of them in the following rule:

(17) Aux → Inflection (M) (have en) (be ing)

Chomsky proposed that the Verb must combine with Aux in all sentences. The brackets around the elements in (17) indicate that it is optional to include the bracketed elements.

As it stands, the phrase structure (PS) rule in (17) will lead to gibberish sentences like *John Inflection may have en be ing kiss Mary*, when Aux is rewritten as "Inflection modal have en be ing." It is clear what the problem is. Some of the elements in (17) are not in their proper place. Although it is true that perfect *-en*

and *have*, and progressive *-ing* and *be*, go together (and this is what (17) expresses), they are not pronounced as a unit; *-en* goes on the verblike thing immediately following *have*, and *-ing* goes on the verblike element immediately following *be*, as depicted in (18).



(18) represents the cross-serial dependencies PS grammars cannot capture nicely. Rewrite rules capture co-occurrence restrictions straightforwardly, but they can't do that and at the same time get the word order right. To remedy this problem, Chomsky argues that PS grammars must be supplemented with transformations, a device that moves elements around (capturing the property of language known as *displacement*).

Taking (19) as a point of departure, transformations will yield (20).

(19) John Inflection may have en be ing kiss Mary

(20) John Inflection may have en be ing kiss Mary

= John may have been kissing Mary

Once available, the device of transformations can be put to use in a variety of contexts. It's in fact what allows linguists to relate structures underlying active and passive sentences, declarative and interrogative sentences, which native speakers intuitively feel are related.

For example, transformations allow one to switch the order of the subject and the first verblike thing and thereby relate *John can swim* and *Can John swim?* Both sentences will have a common structural core (captured by rewrite rules), and transformations will account for the different word orders. Likewise, by adding an auxiliary and switching the subject and the object of a sentence, one can straightforwardly relate the sentences *John arrested Peter* and *Peter was arrested by John*.

To sum up this discussion, Chomsky established in the earliest days of the modern cognitive revolution that a fair amount of abstract structure is necessary to capture some of the most basic properties of what we (tacitly) know about our language. Specifically, the property of infinity that Descartes had already highlighted requires a recursive mechanism (a mechanism that makes precise Humboldt's demand for "infinite use of finite means") that goes beyond what could be achieved if words were like beads on a string (which our perception of sentences as linear strings would lead us to believe). The right mechanism is one that constructs invisible hierarchies, establishing structural relations which can then be manipulated by transformations to fully capture the associations among words that our language faculty spontaneously computes. Without such mechanisms, our model of mental chemistry¹⁸ for

language would be worthless. With them, we have at least a chance of looking in the right direction, equipped with the right basic (mental) structures.

3 More (Hidden) Properties

Needless to say, this is just the beginning. There is more, much more, hidden in the fabric of language. So far we only have the bare minimum set of properties: discreteness, infinity, hierarchy, displacement. Now I would like to spend some time on additional properties that to this day capture linguists' attention, and demand a rich model of the mind.

3.1 Fractality

The first one on my list is often called uniformity, but I prefer to call it fractality. Fractals¹⁹ are those mathematical objects that multiply themselves at smaller scales, a bit like the Russian dolls I mentioned above. Consider the triangle in Figure 4.3a. Now, insert a smaller triangle inside it, and flip it upside down, as in Figure 4.3b. Notice how you have thereby obtained three smaller triangles that are identical to one another and that look like smaller versions of the big triangle in Figure 4.3a. Now take any of these three triangles and repeat the action as in Figure 4.3c.




Figure 4.3 Fractalizing a triangle

Again, what you get is three smaller triangles that are identical to one another, and that look like the bigger triangle that got “cut.” You can repeat this operation ad infinitum. At some point you won’t be able to perceive the small triangles, but they will be there. Figures obtained by fractalization can be the source of endless fascination, and have been used (without awareness of the notion of fractal, which is pretty recent) for all sorts of artistic and religious purposes.²⁰

Consider the beautiful Figure 4.4 on p. 70, obtained after several applications of the rule discussed above. Furthermore many natural objects are organized in a fractal fashion. If you have ever peeled an onion, you know that you end up with concentric layers, one inside the other. If you cut a head of broccoli and look close enough, the bits of broccoli you obtain look like small versions of the broccoli you started off with. Nervous systems and arteries likewise exhibit fractal properties




Figure 4.4 Triangle showing fractalization




Figure 4.5 The human nervous system and arteries


(Figure 4.5 on p. 70). So, when I say that linguistic structures are fractal, they are certainly in good company. But in a cognitive context their fractality is more than mere curiosity. It is at the heart of a fundamental property in language known as “headedness,” “endocentricity,” or the “periscope” property.²¹

3.2 Headed hierarchies

I have already mentioned that in between the macromolecule of our mental chemistry we call a sentence and the atoms we call words, there exist units of computations known as constituents or phrases. We already saw a few of them when we discussed rewrite rules. In the fragment of simplified English that I used, there were noun phrases (NPs) and verb phrases (VPs). Now, there are many more phrases than these two in natural languages;²² exactly how many is still a matter of debate. But the interesting thing for us is that all the phrases that we are sure about, and, in fact, all those we anticipate, have the same organization. Phrases are like the physicist’s electron:²³ once you’ve seen one, you’ve seen them all.

You may have thought that a noun phrase would be different from a verb phrase, but really deep down, they are just two instances of the same general phrasal molecules, like two sides of the same coin. The representational invariance across phrases was established by Chomsky in the late 1960s.²⁴ It was the similarity he noticed between *The enemy’s destruction of the city* and *The enemy destroyed the city*, or between *The city’s destruction by the enemy* and *The city was destroyed by the enemy* that led him to claim that noun phrases and sentences really have the same structural representation.²⁵

The structural representation Chomsky assigned to both sentences (VPs for him) and NPs was roughly as in (21) below:




Chomsky further observed that nominal groups like [*John’s portrait of Saskia*] and verbal groups like [*John portrayed Saskia*] were not all that different from adjectival groups like [*very fond of Saskia*] and prepositional groups like [*right above Saskia*].


In each case, Chomsky noticed, there is a key element which gives its identity to the whole phrase: the noun for noun phrases, the verb for verb phrases, the adjective for adjectival phrases, and the preposition for prepositional phrases. Chomsky called this key element the “head” and the identity of the whole phrase the “label” of the phrase. The head can be modified by material to its left and to its right. Chomsky noted that the material to the left of the head in English can be more easily omitted than the material to the right of the head. For example,

most prepositions need to be followed by some material (**for* __; **of* __; etc.), but none need to be modified by elements like *right*, as in (*right*) *above* Saskia. For this reason Chomsky claimed that the material to the right of the head, which he called the complement, is more tightly connected to the head, and that the head and complement form a subunit (a subphrase, if you want) to the exclusion of the material to the left of the head, which Chomsky called the specifier. Chomsky called the subunit consisting of the head and the complement the “bar-level” category (in writings about syntax you’ll see the bar-level category written A' or \bar{A} , in both cases pronounced “A bar”). Putting all this information together, the structures for noun phrases (NPs), verb phrases (VPs), adjectival phrases (APs), and prepositional phrases (PPs) are given in (22):


(22) a.




b.



c.




d.




Chomsky hypothesized that all four phrases are the reflexes of a more abstract category he called X (a variable varying over $\{N, V, A, P\}$) and represented as follows (the so-called X -bar schema):

(23)



According to the X -bar schema, sentences are but giant phrases, which, once you peel them, reveal smaller but identical phrases, which also contain identical phrases inside them, and so on until you hit upon words (morphemes, really) – atoms you can’t break down any further.



Of all the characteristics of phrases, the existence of an element inside them that's singled out as the head is quite remarkable. It gives linguistic hierarchies their special flavor, for not all hierarchies that may look like linguistic tree diagrams pattern this way. In fact, some linguists think that headedness is perhaps unique to language among cognitive domains that represent information hierarchically.²⁶ Think of family trees (and the way we mentally structure social relations). There is nothing about them that forces one element of each constituent to be more prominent than the others. But somehow when two words combine linguistically, one of them “wins.”

It is this property that captures the periscope effect that is introduced in every textbook dealing with syntactic constituency.²⁷ There are a few tests for constituency that linguists have come up with. One of them is displacement: you can only manipulate phrases. For example, starting with a structure corresponding to something like *John bought Mary's book*, you can say *It is Mary's book that John bought*, but not *It is Mary's that John bought book*. That's because *Mary's book* (not just *Mary's*) is the whole (noun) phrase. Another test for constituency is “replacement.” You can often replace whole phrases (but not parts of them) with just one word. For example, you can replace *the girl next door* with the pronoun *she*. Crucially, in the context of our discussion of the periscope property, the little word that replaces the whole phrase must be of the same type as the head of the big phrase it replaces. So a noun phrase can only be replaced by a noun, a verb phrase by a verb, etc. By this test, you know that a phrase like *The Paris that I visited when I was a boy* (*was a gorgeous city*) is a phrase whose head is *Paris* because you can say “*Paris was a gorgeous city*” (but you can't say “*the was a gorgeous city*” or “*that I visited was a gorgeous city*”). It's as if phrases were built from within, with the properties of the head forming an umbrella that encompasses the space occupied by the whole phrase.

Such properties as the periscope property are what linguists have in mind when they say that linguistic computations are structure-dependent,²⁸ meaning that the units that are being manipulated depend on structural properties like headedness, hierarchical organization, etc.

3.3 From hierarchy to linear order

Headedness also comes in handy for another important feature of linguistic structures. If indeed words are not like beads on a string, the hierarchies that we

describe in terms of tree diagrams must somehow be mentally related to the linear strings we perceive and produce. Put more technically, the hierarchical structures that are used to capture fundamental properties of languages must be turned into (technically, mapped onto) linear structures corresponding to the objects of speech perception and speech production. How is this done? The best hypothesis we have is what's called a linearization procedure.²⁹ Linearization is a process of dimension reduction. It takes a two-dimensional linguistic structure and collapses it onto a single line, as shown in Figure 4.6.




Figure 4.6 Linearization

Mathematicians (and physicists) have worried about dimensions and how to get from higher to lower dimensions for a long time.³⁰ There is a whole field, called projective geometry,³¹ that focuses on such issues. It's an issue painters have to wrestle with, too, for they have to translate 3-D objects onto a 2-D canvas. It's a problem we constantly encounter (though unconsciously) in the opposite direction, when we have to recreate a 3-D shape from a 2-D signal on our retina (linguistically, we also have to recreate a 2-D representation from a 1-D string when we receive a message). So central is this problem that the great vision scientist David Marr devoted much of his energy in his landmark book³² *Vision* to it. It's as central as hierarchy in linguistics. In Zygmund Pizlo's words,³³ it's "the" problem in machine vision.

Dimension-reduction is not always transparent. In fact, mathematicians tell us that there is always some loss of information as we move from higher to lower dimensions.³⁴ In the linguistic domain, this is what gives rise to "structural" ambiguity: at which hierarchical level do we "attach" *with binoculars* in *I saw the man with binoculars*: do we associate it with *man* (man with binoculars) or with *saw* (saw with binoculars)?

Linearization is one of the hottest topics in linguistics currently. There are still many steps we do not understand about how we get from a syntactic structure to a linear structure. But most linearization algorithms currently on the market³⁵ make use of the asymmetry between head and non-head within each phrase to somehow order words on a one-dimensional line. That is, most algorithms have the effect of translating "higher than" into either "precede" or "follow"; so if "x is higher than y, x (say,) precedes y." In other words, headedness becomes the anchor point for the linearization process. Without the hierarchical asymmetry coded via headedness, the (mental) machine responsible for linearization would not know where to start.

There are two more aspects of mental chemistry I would like to briefly touch on before moving on to some other domains of research. One is locality; the other, conservativity.

3.4 Locality

Locality³⁶ refers to the fact that although the portion of our cognitive apparatus dedicated to language can compute potentially unbounded dependencies of the “nested” and “crossing” kinds, there are some restrictions as to which dependencies are legitimate. Consider the following sentence:³⁷

(25) When did John say he fell?

Part of your linguistic knowledge amounts to being able to understand the sentence in (25) as being ambiguous. It can be a question about when the saying took place, or it can be a question about when the falling took place. Similarly for (26) (although one of the two possible questions “How did he say it?” is not very natural, but it’s possible; you can imagine asking about whether John said it in a loud voice, or whispering, etc.):

(26) How did John say he fell?

Now, one interesting property of natural languages is that they allow you to ask several questions at a time (think of this: *Tell me: who did what to whom?*), but they impose restrictions on what exactly you can ask in one sentence, and how you can ask about it. Thus, you can collapse the question in (25) and the question in (26), as in *When did John say how he fell?*, but notice now that this can no longer be a question about when the falling took place. It can only be about when John said that he fell (somehow). In other words, the sentence ceases to be ambiguous.

Why is that? The best answer we have is that the presence of *how* somehow blocks the establishment of a dependency (displacement relation) between *when* and *fell*. In the absence of *how* (as in (25)), *when* and *fell* can be related, but the lack of ambiguity in *When did John say how he fell?* seems to indicate that *fell* can only associate with the question word that’s closest to it. *When* is now too far from it. It’s a bit like when a magnet ceases to attract a nail because you put something in between them that causes the magnetic field to change. Physicists in such cases talk of a locality effect, and linguists do the same. In fact, linguists have reached a remarkable conclusion in this domain: Somehow, linguistic principles conspire to always yield the shortest dependency possible.³⁸ It’s because you can establish a shorter dependency between *fell* and *how* in our example that you cease to be able to establish the dependency between *fell* and *when*.

Here too there is an interesting parallelism with vision. Many vision scientists, since the Gestalt psychologists almost a century ago (modern cognitive scientists *avant la lettre*),³⁹ proposed principles of cognitive shape formation that demand that the “simplest” form be constructed whenever a choice arises. For example, if you have the impression that a white strip is on top of a black strip, in something like Figure 4.7 on p. 76, your visual system assumes that the part of the black strip that’s occluded has the simplest conceivable shape (straight lines). So simplicity




Figure 4.7 White strip on black strip

(“as straight as possible” or “as short as possible”) seems to be a rule active in both visual and linguistic interpretation.

The reader may recall that in Chapter 1 I briefly touched on a parallelism requirement in both the linguistic and visual domains that is also part of how we interpret cognitive objects. Thus, the requirement that we must interpret conjoined objects in parallel makes it impossible to manipulate one object and not the other,⁴⁰ which is why you can’t ask a question about one object, but not the other in:

- (27) What did John eat cheese and? or
 What did John eat and cheese?

Although you can perfectly well ask:

- (28) What did John eat? or
 What did John eat with cheese?

3.5 Conservativity

The final property of linguistic computation I would like to address is called conservativity, and it’s intimately related to the property of displacement.

Even if *ing* and *be* are not close to one another in *John may be swimming*, your cognitive system for language system makes it possible for you to treat them as a unit (indeed, it requires you to treat them as a unit), as we saw above. Likewise, even though *what* and *bought* can be arbitrarily far apart in *What did John say (that Mary believed that Sue said that . . .) Bill bought?* you have no problem understanding that *what* is related to *bought* – that’s why the following sentence will leave you puzzled:

- (29) What did John say that Bill bought cheese?

Similarly, when I say *It’s pictures of himself that Mary said that Bill destroyed*, you know that the word *himself* refers to *Bill*. In all such cases your language faculty

tells you that the relevant element (*ing*, *what*, *pictures of himself*) has been displaced, and ensures that the relation that exists between *ing* and *be* or *what* and *bought* in my examples is conserved no matter how far away from one another these elements end up. It's as if, linguists hypothesized, there was a silent, invisible "trace"⁴¹ or "copy"⁴² of the element that's been displaced – a footprint, as it were. It's as if *what* were in two positions at the same time: at the beginning of the sentence and right next to *bought* (and that's why you can't pronounce anything else, like another object, in that position). Once established, linguistic relations in a sentence can be distorted, but never broken. (In this sense, studying linguistic relations, especially those involving displacement, is a bit like doing topology,⁴³ a subfield of mathematics concerned with the equivalence among shapes that can be deformed into one another without cutting or gluing; e.g., a coffee cup with a handle and a doughnut.)

The silent trace/copy is the device linguists use to code this fact. I close this section by showing you that devices like this one can shed light on some facts one would ordinarily never think twice about. Take the fact that in colloquial speech, a sequence like *want to* is often shortened to *wanna*, as in *I want to go to the movies* → *I wanna go to the movies*.⁴⁴ Now consider the following sentence:

(30) Who does the coach want to shoot?

This sentence is ambiguous; it can mean 'Who is it that the coach wants to do the shooting?' or 'Who is it that the coach wants dead?' Now consider the following:

(31) Who does the coach wanna shoot?

All of a sudden, the ambiguity is lost; this sentence can only mean "Who is it that the coach wants dead?"

Linguists have a nice explanation for this fact: if displaced elements leave a silent trace behind, there are two possible spots *who* could have left a footprint in: between *want* and *to* (the position occupied by *Bill* in *the coach wants Bill to shoot*) or after *shoot* (the position occupied by *Bill* in *the coach wants to shoot Bill*). But once *want* and *to* are fused into *wanna*, there is no more room in between the two of them for the footprint of *who*. Accordingly, the only possible footprint location is after *shoot* and the sentence *Who does the coach wanna shoot?* is indeed predicted to be unambiguous.

It's little gems like these that reinforce a linguist's opinion that the formalism is on the right track.

4 A Note on Idealization

Let us take stock. We have discussed several properties of our cognitive system dedicated to language: discreteness, infinity, hierarchy, endocentricity, locality,

conservativity, and linearity. All of these properties are universal; they are part of what is needed to describe any specific language. As we saw, all of them require a fair amount of abstraction and formal description. In my experience this is the pill that's hardest to swallow for students. They somehow acknowledge that facts like "wanna-contraction" must be explained, but they find it hard to accept that these facts require explanations in terms of rewrite rules, transformations, traces, and the like. "Could all this really not be learned, somehow?" The answer is, and indeed must be, a resounding no.

To me, this inclination to invoke culture, as opposed to nature, is a source of endless puzzlement. Students taking a course in physics, or math, or molecular biology, expect the material to be technically challenging, and they come to understand fairly quickly that explanations that "feel" right virtually always turn out to be seriously mistaken, to be replaced with notions that are hard to grasp and concepts that can't be seen or touched. No such expectation is part of what the first-time linguistics student anticipates, and it is therefore important to stress that cognitive science is just like any other science. It's technically challenging, and far from intuitive. It too seeks to render complex visible phenomena intelligible by reducing them to a few simple, invisible elements.⁴⁵

It may be worth mentioning at this point that the urge to simplify (in service of the urge to understand) is manifest at various levels of analysis: data being collected, questions being raised, and theories being proposed.⁴⁶

No linguist/cognitive scientist seeks to capture linguistic facts one could collect by taping conversations, dissecting news reports, etc. That would be like asking a physicist to tape events that happen in the world, and explain them. That's hopeless. Gathering data is part of the experiment. You try to isolate a few variables, and hope to control for all the irrelevant factors that bear on a few simple questions that you hope are answerable. Science is the art of the soluble. Science is very humbling in this sense. In the jargon of philosophy of science, the key process is called idealization. And idealization is as central to cognitive science as it is to the better-established natural sciences. Think of the toy language *Martian2*; it only has two words (*a* and *b*) to which I haven't even attached any meaning, but it allows us to see the limits of finite-state machines very clearly. Think of the property of infinity at the core of language. This is not something that can be experienced directly. No one will ever produce an infinitely long sentence. But we would be missing something big if we didn't make a distinction between what Chomsky has called competence and performance.⁴⁷ Competence is a synonym for what I have called knowledge of language. Performance refers to specific speech acts during which that knowledge is put to use. Everybody knows that such acts will only partially reflect the knowledge underlying them, and each utterance will be the result of a myriad of factors on top of competence: it will involve the speaker's state of mind, the peculiar context of use, the memory limitations, the zillions of thoughts taking place in parallel, the false starts due to stress, and so on. Trying to account for all of these things would make our theory of mental grammar much more complex than required – in fact, it would make formulating such a theory impossible. Imagine

what a theory would look like if it had to predict that the longest sentence on record contains, say, 1,675 words! What would the rules have to look like to make these facts true, and what would happen if tomorrow someone comes up with a sentence of 1,676 words? The theory would likely have to be rebuilt from scratch.

So, the property of infinity is a simplifying assumption.⁴⁸ It's an idealization (on a par with Turing's "infinite tape" discussed in Chapter 2), one that allows for simple questions to be asked, and hopefully simple theories to be constructed.

As you, as it were, purify the data, you begin to construct simple models that capture the essence or the core logical properties of the system you are making hypotheses about, knowing full well that such models are but very preliminary steps toward the truth. This process amounts to looking for the fundamental abstractions that would shed light on a few questions. As the biochemist Albert von Szent-Györgi⁴⁹ put it, "discovery consists of seeing what everybody has seen and thinking what nobody has thought." Most important discoveries have been the results of puzzling over familiar facts, things people take for granted. Einstein loved to say that this was the result of relinquishing that child-like curiosity later than most adults do.

5 Marr's Three Levels of Analysis

The great vision scientist David Marr⁵⁰ made it clear that the first task to focus on for the cognitive scientist is what he called the "computational" aspect of the problem. Make sure you first capture the logical properties that the system (brain) must have to account for the most basic of behaviors. I hope to have shown in this chapter that Chomsky's results in *Syntactic Structures* took a giant leap in this direction. I think it's fair to say that in the case of many cognitive capacities we haven't reached that stage yet. The second task, according to Marr, concentrates on what he called the algorithmic level. Try to find the best (simplest, etc.) mechanism that would compute the functions that must be computed and that have been identified at the computational level.

Currently, linguists think that the best way to think of what our knowledge of language amounts to consists of a subtle interplay of a few basic operations:⁵¹ a process of lexicalization (turning concepts into words); a core recursive operation, called *Merge*, that takes words two by two and combines them into a set (take word 1 and word 2 and combine them into a set, take that set and combine it with word 3, take that set and . . . *ad infinitum* – with the possibility of recombining word 1 and the entire set that contains it); an operation marking one of members of the Merge-set as more prominent than the other (the way some contours or edges are more prominent in the visual field, which could relate to the issue of endocentricity); and, finally, a way to segment this sequence of set-formation into smaller groups, not too dissimilar from what George Miller⁵² hypothesized for working memory in a seminal paper at the very beginning of the cognitive revolution. Miller observed

that we have an uncanny ability to break down long sequences we have to remember, like phone numbers, or shopping lists. Miller noted that the content of what must be kept in memory does not matter. We seem to break sequences down in the same way; we “chunk” them into smaller parts, with a maximal size of 7 (+/-2) for each chunk. In other words, we spontaneously organize sequences into subhierarchies, like this:

$$(32) \quad \begin{array}{ccccccc} & & (X & & X & & X) \\ & & \underbrace{\hspace{1.5cm}} & & & & \\ (& X & X & X &) \\ \underbrace{\hspace{0.5cm}} & \underbrace{\hspace{0.5cm}} & \underbrace{\hspace{0.5cm}} & & \\ (xxx) & (xxx) & (xxx) \end{array}$$

Chunking is a superbly efficient space-saving device that also limits the domain over which dependencies can be established. At the moment, the best way to understand the locality constraints mentioned above in the context of language seems to be in terms of how we chunk Merge-sequences, and which edges can be detected in these chunks.

Needless to say, the identity of the basic operations at the heart of the language faculty is a matter of intense investigation, with lots of specific proposals, problems, and alternatives to sort through, and I won't be able to do justice to them here. What I want the reader to realize is the characterization of Marr's computational level of analysis acts as a boundary condition for hypotheses at the algorithmic level. The algorithm should be as simple as possible, but not too simple, for it must compute all the things that we know we must be able to do to account for what we (tacitly) know about our language.

The third task, according to Marr, deals with neural implementation. The (aspects of the) knowledge I described, the computational properties of cognitive faculties like language, are properties of the brain, specifically the neurons and their firing patterns. Exactly how the properties I described emerge from this most complex organ is not known. Suffice it to say that we do not currently understand very well at all the neurophysiological substrate of our mental structures. Some see this as a very unfortunate state of affairs, one that threatens the legitimacy of cognitive science, but this is an overreaction. We seem to be in the same boat chemists were in the middle of the nineteenth century, when many aspects of chemical compounds were rendered intelligible in terms of abstract notions like valence, even though the physical basis of valence was not properly explained until many decades later.⁵³ The lack of a physical explanation for valence did not invalidate the chemists' work – in fact, the results of chemistry were interpreted as boundary conditions to be met by the new physics. A similar conclusion should be put forward in the context of the relationship between mental structures and neurophysiological substrate. What cognitive scientists discover now is what the neuroscience of the future will have to capture. In the absence of proper characterization in mental terms, how would the neuroscientist know what to look for? As progress is being made on the neuroscience front, the best cognitive scientists can do is establish a

body of doctrine of mental chemistry, to be used as the refined goals of neurophysics and biology. The past 50 years of cognitive science have demonstrated that considerable progress can be made even in the absence of a clear picture of how cognition arises from neurophysiology. We can only hope that the neuroscience will catch up, for Marr's three-step approach will be most effective when the three levels of analysis are able to constrain one another: some algorithms will have to be abandoned if they are provably beyond the reach of what the neuroscientists deem possible (for a brain like ours), refinements will have to be made at the computational level, or neuroscientists will be asked to reconsider their findings. If such a state of understanding ever arises, the field of cognitive science will be able to rely on data from three distinct domains – mind, brain, and behavior – to support its claim, and will be able to use converging evidence (with emphasis on *converging*: data from behavior will be no less suggestive than data from the brain) to solidify its results, and make them falsifiable in more ways than one, as any science should do.

But until that day comes, the explanatory power of mental structures must not be underestimated. To date it is the most explanatory level of analysis we have, and we should therefore feel very confident in our pursuit of computational properties of the sort discussed in this chapter.

5

The Variety of Linguistic Experience: The Towers of Babel and Pisa

1 The Paradox

To the student of language, the diversity of linguistic systems is very striking. I'm sure you have experienced at least once the helplessness that settles in when you are surrounded by people who speak a different language. Current estimates reveal that there are some 7,000 different languages being spoken on the planet.¹ I'm not sure how this figure was calculated, but I know it's false, for I don't know of any non-arbitrary way of calling something a language and another thing a dialect.² For linguists, both languages and dialects are linguistic systems with equally interesting things to tell us.³ If you factor in the number of dialects, the number of linguistic variants explodes. And, since we are taking an I-linguistic perspective in this book, with the emphasis on the *Individual*, a good case can be made that there are as many distinct linguistic systems as there are individuals on the planet.⁴ For sure, some variants are virtually indistinguishable, and others are as far apart as you could imagine, but one thing is certain: we are dealing with a very large set of possible languages. Moreover, we shouldn't forget that the mental structures underlying our knowledge of language must account not only for the existing (actual) languages, but also for the languages spoken by our ancestors, and all the languages that will be spoken in the future. In short, it must account for every *Homo sapiens*. We have no evidence that the linguistic computations underlying dead languages were any different from those assumed to capture languages currently spoken. The ancient languages that we know of were not more or less complex than the ones we have now; and the new ones that are emerging (recall Nicaraguan Sign Language in Chapter 3) are just as complex as the ones we are more familiar with. It's the same cognitive capacity underlying them all.

So, we quickly face the well-known paradox: how can the same cognitive system give rise to the bewildering complexity that both fascinates and frustrates us?

To make this paradox more vivid, the linguist Mark Baker has brought up the story of the Navajo Code Talkers.⁵ During World War II, the Americans had the brilliant idea to use the best natural secret code available to counter the Japanese: the Marines recruited Navajo speakers to transmit messages. All it took was one bilingual Navajo–English speaker coding the English message into his native language, and another bilingual Navajo–English speaker listening to it and translating it back into English on the other end of the communication line. The Japanese, who had been successful at breaking previous codes, failed miserably at this one. And herein lies the paradox: On the one hand Navajo must be very similar to English for the translation to succeed and the message to be conveyed without distortion. On the other hand, Navajo must be very different from English to be able to confuse the Japanese code-breakers, who could deal with English just fine. Two languages, very similar, and at the same time so distinct. How can we make sense of this?

2 Parameters

It's fair to say that until the late 1970s,⁶ barely 30 years ago, the paradox was a complete mystery. Somehow linguists knew that the same language faculty must underlie the rules of English and Navajo (and Japanese, and all the other possible languages), for this is what the acquisition process tells us. No child finds it hard to learn English or Navajo if placed in appropriate circumstances. As generations of immigrants have been able to witness, language development has nothing to do with the language your parents speak. Placed in the right environment, with the right amount of input, any child can learn any human language. All the languages must be the same, somehow. But if they all must be the same, why do they feel, sound, look so different? And why are they so hard to learn after the onset of puberty?

The breakthrough came at a conference that took place in Pisa in 1979.⁷ Right there and then Chomsky managed to conceptualize language variation in a radically new way,⁸ a way that removed at once the impression of a paradox, and that has since then be part of our standard model of the language faculty. It was one of those Eureka! moments that, once they happen, make you wonder why no one thought of this before, because it makes so much sense (this is what the great physicist Richard Feynman reportedly took to be the telltale sign of a right idea).

The novel idea Chomsky introduced is the notion of a parameter.⁹ He asked us to imagine Universal Grammar as a system of rules, many of which can be switched on or off, or have a variable that needs to be set from a list of possible values, and he further asked us to imagine the process of language acquisition as an eliminative process – a way of removing all but the one possibility that gives the correct output for the language spoken in a child's environment (or, I should say, the

language that the child thinks is spoken in her environment, for we will see in the next chapter that sometimes this can make a big difference, and can result in the development of a new language).¹⁰

Recall from Part I that Universal Grammar can be thought of as a set of *don't's*, more than a set of *do's*. The rules of Universal Grammar tell you what not to do – in any language; in other words, they tell you what to do, but indirectly. Now Chomsky is asking us to add to each rule a list of possible *do's* to choose from. For example, say, Universal Grammar tells you that in order to form an interrogative sentence, there must be a special question word like *who*, *what*, *where*, etc. (“Don’t form a question without a question word.”) But stated as such, the rule doesn’t tell you where the question word may be placed: right next to the verb it modifies, at the beginning of the sentence, etc. Parameters can be thought of as a menu list for each rule, giving you all the options you have (excluding quite a few logical possibilities, too). The learner’s task is to pick the right one (or, in multilingual environments, the right ones).

According to the parametric approach, we can conceive of Universal Grammar as a rule-network akin to a circuit with a set of on/off switches.¹¹ The specific pattern of switches (on for rule 1, on for rule 2, off for rule 3, etc.) yields a specific mental grammar, corresponding to a specific language variant, like the schema in Table 5.1.

Table 5.1 Parameter switchboard

<i>UG</i>	<i>Rule 1</i>	<i>Rule 2</i>	<i>Rule 3</i>	<i>Rule 4</i>	<i>Rule 5</i>	<i>Rule 6</i>
English	ON	OFF	OFF	ON	ON	OFF
Chinese	OFF	OFF	ON	OFF	ON	OFF
Russian	ON	ON	OFF	OFF	OFF	ON

Parameters are actually more subtle than statements like “language A has rule 36, and language B doesn’t,” which the switch-board metaphor might lead one to imagine. Both language A and language B (along with all the other possible languages) have rule 36, but parameters tell you to put one specific element involved in rule 36 here or there, or how much of it here and how much of it there, or to put this element here or another element there.

3 A Few Examples

A few concrete examples of parameters may help here.¹² Consider the fact that in English questions that seek more than a yes/no answer must *start* with a question word, as in *Who did you see?*, or *What did you buy?*. In Chinese, no such requirement exists regarding the position of the interrogative word.¹³ The question word shows up in the same position as the word providing the answer does:

- (1) Ni mai-le shenme?
 You buy-past what
 'What did you buy?'
 (2) Ni mai-le yiben shu.
 You buy-past one book
 'You bought a book.'

Once we remember the *wanna* facts mentioned in the previous chapter, specifically, the fact that rules of grammar exhibit the property of conservativity, the difference between Chinese questions and English questions could (and indeed has been argued to)¹⁴ reduce to where to pronounce the question word: in its displaced position, as in English, or in its non-displaced position (as in Chinese). So the parameter could be formulated as this: "Pronounce question word at the beginning of the sentence?" Chinese learners say no, English ones say yes. A simple on/off switch.

Likewise many languages such as Italian or Spanish allow the subject of sentences to be omitted, as in: *amo Maria* 'love Maria.' English typically requires the subject to be overt. Again the difference seems to be a matter of pronunciation.¹⁵ As the following table shows, Italian (and languages behaving like it) places a rich amount of information in the inflection of the verb.


io	amo	I love
tu	ami	You love
lui, lei	ama	He/She loves
noi	amiamo	We love
voi	amate	You (plural) love
loro	amino	They love

It's no surprise that the subject can be omitted in such languages, as the inflection on the verb makes it clear who the subject is. English can't do this, as it only inflects verbs in the third person singular (loves); the only other verb form (*love*) is too unrevealing. So here, the amount of information leads to a difference that again reduces to where things get pronounced.

Consider now the placement of adverbs in English and French.¹⁶ In English, an adverb like *quickly* may not intervene between the verb and the direct object, in contrast with French.


- (3) a. *John eats quickly an apple.
 b. Jean mange rapidement une pomme.
 c. John quickly eats an apple.
 d. *Jean rapidement mange une pomme.

Suppose that in both (in fact, all) languages, the clause has a structure roughly as in (4), with Infl short for "Inflection."¹⁷




Assume also that the Verb (V) and the inflection (Infl) must combine in both (/all) languages. Based on (4), one can see that the resulting association between V and Infl could be pronounced where Infl is, or where V is. This has been argued to be the relevant parameter for adverb placement. If the verb is pronounced at the Infl position, we get French; if it's pronounced in its V position, we get English. Once again, the difference reduces to an on/off switch.

(5) a. English



b. French



As a final example of the logic of parameterization, consider the fact that in languages like Japanese the verb typically comes last in the sentence. Thus, a sentence like *Taro ate an apple* comes out as (6):


(6) Taro-ga ringo-o tabeta. (literally, 'Taro apple ate.')

Experts in comparative linguistics like Joseph Greenberg¹⁸ had long observed that this difference between OV and VO correlates with other differences between the two languages. Just like Japanese objects precede verbs, objects of prepositions precede prepositions (in effect, turning the latter into *postpositions*). So, in Japanese, *with Taro* is *Taro-to* (literally, 'Taro-with'). Likewise, objects of nouns (*portrait of Taro*) precede nouns in Japanese: *Taro-no syasin* (literally, 'Taro-of portrait'). And similarly for objects of adjectives: *kind to Taro* is *Taro-ni sinsetu(-da)* (literally, 'Taro-to kind').


Based on our discussion of the invariance of the linguistic molecule in the previous chapter, I am sure the reader has already thought of the parameter¹⁹ that would

capture all these facts in one single stroke. Recall that the rule Universal Grammar imposes is that there be a head in each phrase. But it does not say where that head should be pronounced: to the left or to the right of the element modifying it? If the language picks the “precede” (head left) option, it will be like English. If the language picks the “follow” (head right) option, it will be like Japanese. Notice that since the headedness requirement applies to all phrases, the parameter applies to all phrases too. What is true of Verb Phrases is true of Noun Phrases, Prepositional Phrases, Adjectival Phrases, etc. So X-bar trees for Japanese in effect end up looking like mirror images of English trees:

(7) a. English



b. Japanese



As research following Chomsky’s seminal proposal has demonstrated, this sort of account can be multiplied to accommodate all sorts of differences across languages.²⁰ Simple parameters, located in the right points in the structure of Universal Grammar, can have massive effects (they can for example make all the phrases look different/inverted), causing the languages to look quite different on the surface, although a single tiny difference is responsible for this.

4 A Lesson from Biology

The proposal Chomsky made next to the famous leaning tower of Pisa can make sense of our impression of mutual unintelligibility among languages, which the Ancients had tried to rationalize in the story of another tower, the tower of Babel.

As far as linguists are concerned, all languages are alike even if they look and sound vastly different. In terms of mental structures, there is basically just one human language, which we may call *Human*. Like Darwin, Chomsky is fond of asking what a Martian would say about differences on Earth. In the case of human languages, if the principles and parameters approach is right, the Martian may well not be able to tell our languages apart, since deep down, at the fundamental cognitive level, they are so similar.

Such a conclusion is likely to upset our common sense, as linguistic differences strike us as fundamental. But consider how you must have felt the first time you were told of another thought experiment performed in Pisa. Galileo allegedly dropped a ten-pound weight and a one-pound weight off the leaning tower of Pisa, and claimed that both fall at the same speed (with the same acceleration). This, too, violates our common-sense intuition. And yet the idealization Galileo

proposed is fundamental to physics. When Galileo asked his contemporaries to disregard friction in the study of motion, he was asking them to remove one of the intuitively most basic facts about motion, but by thinking like no one had thought before, and removing an interfering factor, he opened the door to modern science. Likewise in his thought experiment about language variation, Chomsky is asking us to disregard one of the most basic facts about language: the fact that words are pronounced in specific places. But if you disregard this fact, languages (dialects, etc.) can be seen as essentially the same, and a whole new perspective on what languages are, and how they are acquired, opens up.

I should perhaps point out that linguistics is not the only field where massive surface differences reduce to tiny points of variation. Biologists have come to endorse a similar view. Beginning with the ground-breaking work of Jacques Monod and François Jacob,²¹ biologists have described the genome as a network of switches, where so-called master genes turn on and off other genes, resulting in genetic expressions that can yield creatures as different as the mouse and the elephant. In light of the more recent discovery of highly conserved (similar across species, that is, evolutionarily stable) portions of the genome, as well as recurrent structural motifs, some biologists think that this is basically the only way of making sense of the variation we find in the biological world.²² It is interesting to note that some developmental biologists have characterized the gradual differentiation of embryos in terms of four basic mechanisms that are strongly reminiscent of the kind of parameters postulated in the linguistic literature and discussed above. Thus, Wallace Arthur²³ lists the following four modes of differentiation:

- heterochrony (different timing of gene expression)
- heterotopy (different location of gene expression)
- heterometry (different amount of gene expression)
- heterotypy (change in the nature of the gene product)

It is interesting to note that both in the realm of linguistics and biology the standard wisdom until recently was that variation was basically endless.²⁴ Today the opposite is the case: both linguists and biologists have come to understand that variation is severely constrained by the cognitive/genetic constraints imposed by the very system that makes language and life possible.

5 How Parameters (Must) Interact

Having shown how cognitive scientists came to think of language variation, I now want to turn to a question that many readers may have been pondering. How many parameters are there? How many different linguistic variants can there be?

The short answer to this question is, we don't know yet. Whereas we are pretty confident about the kind of parameters there can be, the number of possible

parameters depends on the number of rules made available by Universal Grammar; and here linguists don't have solid understanding yet. But if recent estimates are anything to go by, there may be as many as 100 parameters. I hope the reader realizes that this is a big number. Assuming (as most linguists do) that parameters are binary choices (on/off), 100 parameters means 2^{100} possible languages (that's 1.27 multiplied by 10^{30} (1 followed by 30 zeros)). Even if we bear in mind that Universal Grammar must define a large set of possible languages to contain all possible linguistic variants (languages, dialects, idiolects . . .), many consider 2^{100} options too big a space for the child to navigate in the course of language development.

How exactly children set the values of parameters is a topic I will address in the next chapter. But here I would like to point out that the huge number we arrived at on the basis of 100 parameters was the result of an assumption that is very likely to be false. The assumption concerns the independence of the parameters. So far I have said nothing about the very real possibility that certain combinations of parametric values cannot even be considered by the child because they would conflict with one another. In other words, it is possible that only one value of a given parameter makes it possible for the child to consider the values of another parameter. To make this more concrete, let us go back to the space of possibilities made available by Universal Grammar in the domain of question formation. The basic parameter discussed above distinguished English and Chinese. But I have said nothing about sentences where more than one question word occurs, as in *Who gave what to whom?* Here we see that English is not too different from Chinese. Although it forces one question word to be displaced to the beginning of the sentence, it prohibits the displacement of more than one question word (*Who what to whom gave?* is not English). This, however, is not a universal prohibition. Slavic languages (like Russian, Bulgarian, etc.) in fact require that all question words cluster at the beginning of the sentence.²⁵ Now clearly the parameter setting apart English and Russian is a choice that can only be considered by children that have first figured out that their language is not of the Chinese type.

The linguist Mark Baker²⁶ has discussed the interdependency of many of the parameters that have been proposed and the literature, and shown how parameters in fact form a hierarchy that dramatically reduces the number of possible combinations of parameter values, hence the number of possible linguistic variants. Figure 5.1 shows the partial hierarchy of parameters proposed by Baker.²⁷ Even though I have not discussed many of these parameters in this book, I hope that this schematic representation helps the reader capture the logic of Baker's proposal. Baker²⁸ goes so far as claiming that the hierarchy of parameters marks the beginning of what one might call the periodic table of parameters, where linguistic variants could be classified with the same predictive power as the table of elements, which defines the space of chemical combinations.

Be that as it may, we have come a long way from the impression shared by virtually all linguists before the cognitive revolution that languages can vary in unlimited and unpredictable ways.

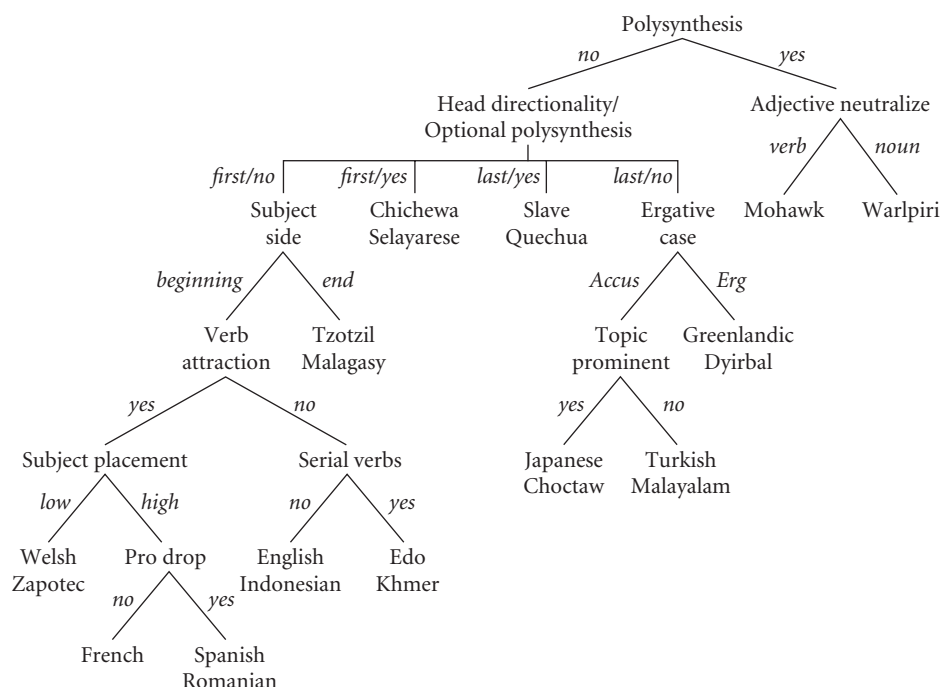


Figure 5.1 Baker's hierarchy of parameters

6 (A Taste of) Comparative Linguistics

Having discussed the notion of parameters, and how it enables the linguist to break through the paradox of unity and diversity, I would now like to focus on one important consequence of the view that all human languages are cut from the same cloth. If this claim is correct, then it vastly enlarges the range of data that the linguist can bring to bear on a specific issue in a specific language. If there is essentially no difference between Swahili and English, data from Swahili can illuminate the structure of English, when various factors internal to English grammar conspire to obscure what Swahili wears on its sleeve. Put another way, the big claim behind the parameter model promotes the comparative method in linguistic analysis. Of course, the comparative method wasn't born with the advent of parameters. Comparing forms and structures across languages was standard in the nineteenth century, in the heyday of philological studies.²⁹ But the purpose of comparison was dramatically different.³⁰ For philologists, comparison was used to argue for a "genetic" relationship (common ancestry) among the languages being compared. It also formed the basis for "reconstructing" the extinct parent language. The purpose of comparison in the context of parameters is much more cognitive in nature. In a parameter model, comparison is used to support claims

about the mental structures made possible by the human language faculty (Universal Grammar).

In the remainder of this chapter, I would like to use an example to illustrate the virtues of comparison in a cognitive context. My example is taken from Morris Halle and Norvin Richards' biographical memoir of the late Ken Hale,³¹ one of the giants of the comparative method in modern linguistics.

Consider the fact that languages throughout the world make a distinction between transitive and intransitive verbs, that is, verbs that take objects (like *kiss*) and verbs that don't (like *arrive*). Many verbs require the presence of an object (think of the fact that you can't say *Bill devoured*), and many verbs prohibit the presence of an object (you can't say *Mary fainted Bill*). Still other verbs allow, but do not require, the presence of an object (*John read (a book)*, *John ate (pasta)*, etc.). Here I will focus on this third class of verbs, the ones that alternate between transitive and intransitive use, specifically verbs like *melt* or *break*:

- (8) a. The chocolate melted.
- b. John melted the chocolate.

- (9) a. The vase broke.
- b. I broke the vase.

From a cognitive perspective, the issue with such alternating verbs is not only why they exist in the first place, but also why some verbs resist such alternations. For example, *laugh* and *cough* don't alternate (indicated by the star at the beginning of the example):

- (10) a. The baby laughed.
- b. *John laughed the baby.
- (11) a. The man coughed.
- b. *I coughed the man.

These facts are quite robust. English speakers don't seem to make mistakes about them: they never produce forms like (10b) or (11b), and grammar classes don't devote too much time to them, because they (correctly) take these facts for granted. And yet, linguists like Ken Hale have demonstrated that the principles underlying these facts reveal the vastness and subtlety of our linguistic knowledge.

If you are like my students, you may be tempted to say that the reason you can't say *John laughed the baby* or *I coughed the man* is because such things wouldn't make sense. But this is far from obvious to me. In fact, I think I know exactly what such things would mean; they would mean pretty much what *John melted the chocolate* and *I broke the vase* mean – something like *x caused y to Verb*: John caused the chocolate to melt, I caused the vase to break, John caused the baby to laugh, and I caused the man to cough. There doesn't seem to be any significant difference

among these expressions; all of them make sense. So there is no problem with thinking the thought that I caused the man to cough; it's just that for some reason I can't express this thought that way. Instead I must say *I made the man cough*. This, by the way, is not a parochial fact about English – it's not the result of a parameter. In language after language, we find the same restriction (including in Navajo, the code-talker language discussed above, which was one of the many languages that Ken Hale was an expert on), suggesting this restriction must be the work of Universal Grammar.

Hale developed a theory of the restriction at issue based on facts from languages like Basque.³² In Basque, the way to express “to laugh” is *barre egin* – literally, ‘laugh do’. The verb for “cry” (another non-alternating verb) is *negar egin*, literally ‘cry do’. So what English expresses as a verb, Basque expresses by means of a generic verb like “do” and a noun that gives the action its flavor. This is not true of all verbs in Basque, but it's true of all verbs that fail to alternate in English.

On the basis of this, Hale hypothesized that the reason verbs like *laugh* and *cough* and *cry* don't alternate in English is because these verbs are to be analyzed at some abstract level in the same way that Basque transparently reveals, as “do (a) laugh,” “do (a) cough,” and “do (a) cry.” Hale pointed out that as soon as you analyze *laugh*, *cough*, and *cry* that way, you have a simple explanation for why they don't alternate: they don't alternate because they are not intransitive verbs (although unlike regular transitive verbs, you don't see the object). Put differently, they cannot be made transitive because they already are transitive; you can't add an object to a verb that already has one (albeit a silent one). Just like you cannot say *John bought the magazine the store*, you can't say *John laughed the baby* because *John laughed the baby* is – suitably analyzed – like *John did a laugh the baby*.

Facts like these reveal the power of the comparative method in the service of our explanatory goals and cognitive aims. At first sight Basque and English are radically different (if you've ever tried to learn Basque, you know what I mean), but deep down, they are really the same; they result from the same mental structures. The differences arise when it comes to expressing these underlying mental structures. In the case at hand, Basque does it more transparently, but no language is privileged in this respect. No language always expresses all its structural foundations transparently; the task of the comparative linguist consists in finding which language(s) express(es) which portion(s) of Universal Grammar more transparently than others. It's very much a collective enterprise.

Of course, once we have arrived at a simple explanation for facts that were not immediately transparent (as in the case of English *laugh*), we must also try to account for the “distortion” if possible; that is, we must try to explain why English does not say “do (a) laugh” – why English doesn't sound like Basque. Quite often this is a question with no interesting answer. It's a bit like asking why the English word for “dog” is *chien* in French. This is what the linguist Ferdinand de Saussure³³ called *l'arbitraire du signe* (the arbitrary character of the signal) – a fact that children simply have to learn by brute force. In some cases the answer may be historical in nature. The reason the English word for “table” is (small details of pronunciation

aside) the same as in French is ultimately due to the Norman invasion of England in 1066. But the reader should be aware that historical answers are not interesting from a cognitive perspective, for at the end of the day, we want to account for how children acquire their language, and children don't have access to historical information of this kind. A child learning the word *table* in English learns that word for no other reason than because that's what the speakers in her environment say. I'm stressing this obvious point because it can be easily forgotten. Very often I hear that the reason English speakers can say *John gave Mary a book* and *John gave a book to Mary*, but not *John donated the library a book* (although they can say *John donated a book to the library*) is because *donate*, unlike *give*, is a word of French/Romance origin, and French does not allow the structure "X Verb Y Z," but it does allow the structure "X Verb Z to Y." This may be true, but this cannot be the piece of information used by the English child to learn that *John donated the library a book* is not part of English.

Having said this, it is sometimes possible to provide more interesting explanations for why certain languages express things differently. In the case of *laugh*, *cough*, etc. Ken Hale hypothesized that the reason these verbs look like simple verbs as opposed to Basque-like combinations of Verb + Noun is because English word-formation has a process of "fusion" that essentially conflates verbs and their nominal objects into one word. Here, too, comparative research can be very instructive. We know that many languages resort to this process of word fusion on a regular basis. So-called polysynthetic languages³⁴ (which include many native American languages, e.g. Mohawk), in fact, regularly, use a single word where English would use a full sentence. Thus, Mohawk uses a single word *wa'kenaktahnimu* to say the equivalent of *I bought a bed*.

In fact, English also shows examples of fusion in various domains. Consider the fact that the verb *broke* reflects "break" + some tense information. This information is fused in *John broke the glass*, but comes apart in *John did not break the glass*, where the auxiliary (*did*) is inflected for tense instead of the verb. Come to think of it, the relation internal to the English language between *broke* and *did break* is quite similar to the relation between the English verb *laugh* and Basque "laugh do": a single verb corresponds to a combination of a generic verb and a more contentful form. It is thus no surprise to see instances of *laugh* being split into two in English, as in *John laughed a really big laugh*. All of these facts strengthen Hale's analysis, which is but one example of the many analyses that linguists have proposed on the basis of comparative research.

Since all human languages are expressions of Universal Grammar, it is logically possible to try to uncover the underlying structure of a language by just looking at that language in isolation, but since specific languages are the result of specific parametric values, we would be missing a lot of the possibilities provided by Universal Grammar by restricting our attention to one language. Comparative work significantly contributes to our attempt to unravel mental structures, and for this reason, the more "linguistic variants" are spoken and studied on the planet, the better our chances of figuring out the nature of our language faculty.

7 Summary

Let me summarize the main points of this chapter. At the end of the 1970s, Chomsky began to conceive of the language faculty as follows. Children come equipped with a set of principles of grammar construction, which we call Universal Grammar. Among the principles or rules of Universal Grammar are some with open parameters. Specific grammars arise once values for these open parameters are specified. Parameter values are determined on the basis of linguistic input the child receives. A particular language's grammar, then, is simply a specification of the values that the principles of Universal Grammar leave open. This acquisition process is sensitive to the details of the environmental input (as well as the level of development of other cognitive capacities), since it is the child's linguistic input that provides the relevant information for setting the parameter values correctly. However, the shape of the knowledge attained (the structure of the acquired grammar) is not limited to information that can be gleaned from the linguistic data since the latter has structure imposed upon it by the rich principles that Universal Grammar makes available.

Much of the work in linguistics since the mid-1970s can be seen, in retrospect, as demonstrating the viability of this conception. At the beginning of the 1980s, there was an explosion of comparative research that exploited this combination of fixed principles and varying parametric values, and that showed that languages, despite apparent surface diversity, could be seen as sharing a common fixed core. The comparative method enables many more facts (from a multitude of languages) to bear on questions that arose in the context of a specific study of a specific language.

Although our linguistic system and visual system are often seen as the two pillars of modern cognitive science, there is a big difference between the two: barring pathology, the visual system is obviously the same for everyone; the language system, by contrast, is perceived as much more varied. But this is our common sense fooling us once again. The linguistic system is as universal as the visual system. However, it is true that our language faculty is sensitive to highly specific properties of the environment, and gives rise to a great number of surface grammatical systems. But in this, the language faculty is not isolated. Our moral sense and our music instinct, for which a strong biological component no doubt exists, also give rise to quite diverse cultural expressions. Perhaps a parametric model exists for them too (as some have recently suggested),³⁵ but, as the history of linguistics makes clear, such a model can only be developed once we have a good sense of the sort of rules (mental structures, computations, etc.) that underlie the relevant capacity. It's only once we have figured out the atoms and basic modes of combinations that we can begin to ask what surface molecules these mental chemistries could give rise to. Research on the moral sense and the music instinct (among many other cognitive faculties) is under way, and it is to be hoped that such research will be guided by the modes of inquiry and results achieved in the language domain.

6

All Roads Lead to Universal Grammar

1 Parameters Are Just a Start

Learning language is easy; there is no learning involved, language is what humans are born with. Yet, at the same time, learning a language is hard, and although the innate component we call Universal Grammar helps tremendously, it still leaves a lot of room for the environment to leave its marks. Children go through the process effortlessly, but this should not lead us to ignore the complexity of the task. The parameter model discussed in the previous chapter cannot provide all the answers related to questions of language acquisition/development. Parameters are a necessary, but not sufficient, component of how we ultimately characterize the learning path followed by every normal child. No one is under the illusion that the parameterization is the whole story – it's the best start linguists ever came up with, but it's still just a start.

Our ultimate theory of language development should not only capture the fact that all children grow a language, and do so in a remarkably uniform fashion, it should also make room for the fact that language learning is very often imperfect, in the sense that children often fail to learn the exact same linguistic variant spoken in the environment – a fact that is responsible for language creation (creolization), and also language change.¹ Something must account for the fact that at one point English children stopped learning Old English and developed a language that specialists call Middle English. Such changes happen frequently (though some are more dramatic than others), and they are an important feature of the language-learning path that must figure in our explanation.

Our theory of language development must also take into account the fact that learning a language takes a while. I'm told² that when the parameter model was first proposed, linguists sympathetic to it started worrying that it made language acquisition too easy: if there are, say, 20 parameters, and learning is just

a matter of adjusting the switches, maybe the child could be done in a matter of days, given the right input. But that's way too fast. Although linguists of a nativist persuasion like myself love to say that language acquisition is very fast, this is like focusing on the glass that's half-full, as opposed to half-empty. Language learning is, from one perspective, quite slow. It takes several years; this too must be accounted for.

Our theory of language development should also be able to capture the fact that although there is a fixed developmental schedule for language, with well-defined stages, there is quite a bit of individual variation: not all children do it in exactly the same way. Language acquisition specialists like Stephen Crain,³ Rosalind Thornton,⁴ and others have provided – to my mind, compelling – evidence that children explore different corners of the space of all possible grammars during acquisition in a way that is not immediately dependent on the obvious fact that no two children receive exactly the same input. So, in this domain too, there is quite a bit of (constrained) freedom and creativity. But as it is said of Rome, all roads ultimately lead (back to) Universal Grammar.⁵

Finally, our theory of language development should wrestle with the fact that any child learning a language must deal with noisy and ambiguous data.⁶ Being exposed to a bit of Chinese every time her parents have dinner at the local Chinese restaurant won't be enough for a child to learn Chinese. We also don't want the occasional visits from Uncle Harry from France to disrupt the learning path of a child focused on acquiring English. In other words, some data count more than others in the course of language acquisition. This is obviously true, but it is not part of the parameter model presented in the previous chapter. Also missing from the parameter model is the fact that children (unlike linguists) don't have a lifetime to figure out which data provide good evidence for one parameter value over another. Take the head parameter. When I introduced it in the previous chapter, I made it look like this: if the object precedes the verb, set the value of the parameter to the "Japanese" value (head-final); if the object follows the verb, set the value of the parameter to the "English" value (head-initial). This makes the task look very, very easy: scan the sentence, figure out where the object is with respect to the verb, and voilà, the right value of the parameter can thereby be set. But this description ignores the fact that all languages show instances of "displacement," where objects are moved away from their original position, as in *What did John eat?*, or *This book, I will never touch again!* In both of these cases, the verb follows the object, so if the English-learning child were not (unconsciously) careful, she could mis-set the value of the relevant parameter. Although it is true that this is likely what happens in the case of language change, we don't want this to happen too often. After all, the default case is for the child to acquire the language spoken in her environment. The bottom line is that the data the child receives are "noisy" and ambiguous, in addition to being fragmentary.

The parameter approach presupposes that the child can deal with this additional difficulty, but it is fair to say that when linguists and other developmental psychologists have sat down to think about how children actually cope with it,

and looked at the sort of input a child typically received, it's proven to be a much harder problem than one might have thought. This is why linguists say that the parameter model shows how the child could *in principle* solve the language acquisition problem; no one thinks that we have found a way to show how it's done in practice. Here we are still in the dark, but a few brave linguists have explored avenues that shed light on the issue,⁷ and that's what I want to cover in this chapter.

To top it all off, the child learning a language must also learn exceptions; she must learn which processes are regular, and which are semi-regular, and which are true of only a handful of words, or perhaps just one. Think about irregular forms (like the past-tense form of *go* being *went*); all languages have some of these, and all children cope remarkably well with them. But how do they do it; how do they segregate the regular from the irregular?

2 More Innateness Is Needed

No doubt in light of the big learning challenge facing them, children throw all their mental resources at it (unconsciously, of course). As Paul Bloom argues in his discussion of how children learn the meaning of words,⁸ the learners use their knowledge of syntax, intonation, social understanding (what's called Theory of Mind, or mindreading, by which we guess people's intentions), and what they have already learned at any given point, to narrow down the space of possibilities to a few, and ultimately, to one. All of these factors, interacting in complex ways, must be involved. As they say, anything helps.

The issue I would like to focus on at this point pertains to how children handle the data. The big distinction I want to talk about is the one between data input and data intake.⁹ Our innate capacities define what counts as data. Without biological priors, everything, meaning nothing, counts as data, as we already discussed in Chapter 2, under the rubric of "information."¹⁰ The bare minimum for learning, the first learning bias, is the hypothesis space – the choices at the learner's disposal; in the context of language, this corresponds to Universal Grammar with all its parameters, and everything else that the child knows (i.e., can learn from). But not every bit of information is equally informative. The data must be distilled. Here the child must be endowed with even more biases, telling her what to ignore, what specifically to focus on at any given point, and so on. Data input (raw data from the senses) is definitely not the same as data intake (what ultimately gets used in the language acquisition process). Finally, given that we are talking about a set of choices, there must also be an algorithm updating the learner's belief in the available hypotheses.¹¹ All of these components of the problem (hypothesis space, data intake, and updating mechanism) can be investigated separately, but at the end of the day, they must all be part of our acquisition theory.

The method of choice used in such investigations is computational modeling. For ethical reasons we cannot directly manipulate the data input, the data intake, and the updating algorithm with real children acquiring real languages. Instead we must construct computer models that simulate the learning task; that is, we must construct simplified scenarios that we hope capture the essential aspects of actual acquisition scenarios. Idealization and simplification are, again, the only path to progress.

The field of computational acquisition is still very much in its infancy, but I think we can already identify a few guiding hypotheses and leading ideas. I will first list them, and then discuss each one in a bit more detail to give the reader a flavor of what they are supposed to do, and how they are supposed to work.

- To deal with the fact that there are regular and irregular processes in language, the child must be equipped with a tolerance principle that prevents her from abandoning a hypothesis as soon as an exception is encountered.¹² It's only when too many exceptions are encountered that the child will be led to adjust her original hypothesis.
- To deal with the fact that various processes like displacement may be compatible with either of two (or more) possibilities the child is considering, or even worse, the fact that such processes may yield outputs that are downright misleading (compatible with what would in the end be the "wrong" choice), the child must be equipped with a data-filtering mechanism that essentially discards ambiguous and other unwanted data points.¹³ This filtering device often takes the form of "cues"¹⁴ – pre-defined (abstract) examples of what counts as incontrovertible, unambiguous evidence in favor of each of the possibilities the child has to consider.
- To deal with noisy data, there must be a pre-specified amount of the data that can serve as enough evidence for a specific value of a parameter.¹⁵ The same amount must also be used to determine when to abandon options that the child has been freely exploring for a while.
- To deal with the fact that it is after all very young children that are involved in this complex learning task, the relevant data (the unambiguous cues) must not be too hard to find, or too hard to keep track of in memory (they must be frequent enough). For this reason we might as well assume that there are certain domains that the child virtually never pays attention to (i.e., never learns from), even if those domains may contain a relevant bit of evidence.¹⁶ The linguist may be able to identify that bit of evidence, but it is psychologically implausible to assume that the child can attend to everything.
- Finally, to deal with the fact that it is a developing human child doing the learning, it is reasonable to assume that not all of the cognitive faculties that the child makes use of are available at the outset of language acquisition; some of these faculties are almost certainly subject to maturational constraints that may slow down the acquisition process.¹⁷

2.1 Statistics

All in all, it seems quite clear that the language-learning process (and no doubt we'd find the same was true of other cognitive domains if we knew enough about them) is channeled by a variety of biases, already at the embryonic stage and throughout the growth period, which restrict not only what the child can ultimately know (the final state of the cognitive capacity), but also how the child can get there. It is also clear that although the cognitive capacity itself is something that you either have or you don't (you can't have 30 percent of the language faculty), the learning process is gradual, and involves probabilistic¹⁸ learning mechanisms that keep track of how much evidence of which kind the organism has been exposed to and has processed at any given stage of development. It is also quite likely that some of these learning mechanisms are not specific to language, but one should not lose sight of the fact that these domain-general mechanisms operate in a domain-specific space of possibilities, defined by our biological endowment; in the case of language, this means Universal Grammar.¹⁹

This point about statistical learning bears emphasis because many recent advances in the field of computational language learning have recognized the importance of this mechanism. Given the fact that much of the data the child receives is noisy, full of ambiguities, and rife with processes that are not 100 percent productive, language development amounts to a gradual selective process during which conflicting evidence forces the child to entertain several possibilities at once, and to weigh the evidence to ultimately decide which option is the winner. Charles Yang is the linguist who has made this point most clearly in his work.²⁰ He has developed a variational model of language acquisition that bears a strong (explicit) resemblance to the process of natural selection known from evolutionary biology. Darwin recognized that population changes can be accounted for by first acknowledging the fact that variation among individuals in the population exists to begin with, which allows for change at the population level to be conceived as a shift in the proportions of individuals in that population bearing the different variants (in so doing Darwin introduced what Ernst Mayr has called "population thinking").²¹

Yang's insight is that the populations that biologists study need not be populations of organisms in order for the logic of natural selection to apply. Yang in fact argues that the mechanisms of selection can apply to a population of languages (each with different traits, i.e., parameter settings) within the mind of the language learner. These languages compete: some get rewarded and others punished as data comes in and matches, or fails to match, their properties (parameters). The probabilities that Sarah will set her parameters for an English-type grammar and not a Chinese-type one keep shifting as data comes in, until learning is complete (that is, until there has been enough unambiguous evidence concerning the identity of the language in the environment). The language the child ultimately acquires is the best fit to the language spoken around her, or, I should say, to the way she has processed the data of the language spoken around her (based on the biases and

filtering devices mentioned above), for our theory must ultimately be responsible for language change and creolization. In fact, the necessity of the learning biases discussed in this chapter makes it clear that the child is not a passive learner. With the help of Universal Grammar, she transforms, filters, and selects aspects of the input, so much so that it is more accurate to speak of language “recreation” than “acquisition” when the language that the child develops successfully “matches” (as far as we can tell) the language of the environment.

Variational models like the one proposed by Yang unify aspects of the two traditionally dominant camps in cognitive science: the “representationalists,” who insist that the mind is a symbol-processing machine with a high degree of specificity to the task at hand (they emphasize that linguistic representations are not like musical representations, and neither are they like arithmetic or visual representations); and the “associationists,” who propose that all behavior can be accounted for by domain-general processes that constantly adjust to the environment. The debate has raged since Descartes and Locke, and as is so often the case with such big divisions, the ultimate resolution of the debate is likely to show that both sides were right in some respects and wrong in others.

But now that modern representationalists have developed a model that can incorporate the favorite tools of the associationists (such as domain-general processes like statistical learning), the modern associationists have been very quick to jump on this apparent concession by the representationalists and to argue for a new wave of empiricism that urges us to “rethink” innateness,²² and question what’s within – effectively reviving Locke’s *tabula rasa*, or the behaviorists’ one-size-fits-all learning process. Let’s be clear: the environment is no substitute for the mind. And statistical learning is no substitute for the domain-specific mental structures that underlie behavior.

As I already made clear in Chapter 2,²³ to be useful, statistics must operate within a well-defined space of possibilities; you can’t get specificity if the hypothesis space is not specific to start off with, so whatever success statistical learning enjoys, it owes it to the pre-specified mental structures over which it operates. One of the clearest examples of the need for highly specific mental structures to regulate domain-general statistical mechanisms comes from the domain that helped launch the new wave of empiricism in cognitive science a decade ago.²⁴

In a widely cited paper, Saffran, Aslin, and Newport²⁵ proposed that one of the central tasks of language learning, picking out words from a continuous stream of speech (“word segmentation”), could be achieved by simply keeping track of transitional probabilities between syllables (an early suggestion of this kind was already made in Chomsky’s 1955 magnum opus, *The Logical Structure of Linguistic Theory*): if syllable B follows syllable A more often than others, then A and B form a hypothetical word. For example, with sufficient exposure, the child may be able to guess that in the four-syllable sequence *prettybaby*, the probabilities of *tty* following *pre* and of *by* following *ba* are both higher than the probability of *ba* following *tty*. Accordingly, the child will hypothesize (correctly) that *pretty* and *baby* are English words, but *ttyba* isn’t. Remarkably, Saffran et al. demonstrate

that 8-month-old infants can in fact extract three-syllable words in the continuous speech of an artificial language after only two minutes of exposure. The Saffran et al. paradigm has been used in a variety of experiments to see how much knowledge could be extracted by means of transitional probabilities alone. But notice that in order to keep track of transitional probabilities in the *pretty baby* example, infants must be equipped with the notion (representation/mental structure) “syllable” – short of which they would not know what to keep track of. Note that syllabification of the continuous speech signal is far from trivial, and in languages like English, syllable size varies: some syllables consist of a consonant and a vowel (*ba*), others consist of a consonant, a vowel and another consonant (*tap*), and others are even bigger (*strings*). The range of possible syllable types in the language must be learned before it can be of any use to compute the relevant probabilities for word segmentation.

Moreover, *all* the possible words in the artificial language used in the original Saffran et al. experiment consisted of three syllables each. But this is clearly not a realistic model for natural languages, where some words consist of one syllable (*bit*), others of two syllables (*habit*), etc. Gambell and Yang²⁶ showed that once this variation is taken into account, the Saffran et al. model fares poorly. This is not surprising since the model is built in such a way that it never posits a word boundary after one syllable (the algorithm looks for some syllable B that follows syllable A more often than any other syllable before it posits a word boundary, but there will never be such a syllable B if syllable A forms a word of its own).

Gambell and Yang go on to show that another innate bias (structural property) may be of great help to cope with this difficulty. It is a well-known fact that no word (in any language) can bear more than one (primary) stress. If we assume (as I think we must) that the learner is equipped with this principle, then she knows that the sequence *bigbadwolf* must contain three words because there are three main stresses in the sequence (bígbádwólf). (This example of the use of prosody information to guide the acquisition task is part of a growing literature on so-called phonological bootstrapping,²⁷ according to which rhythmic properties of the language are among the earliest, if not the earliest, clues the learner uses to figure out properties of the ambient language.)

The Gambell and Yang result demonstrates how crucial innate mental structuring principles are to navigating the learning path, and shows how much the action of domain-general mechanisms depends on domain-specific boundary conditions.

Having stressed this fundamental point, let me now return to some of the additional constraints that are needed for learning to succeed.

2.2 (Ir)regularities

Children are known to overgeneralize, producing forms like *goed* where they should really say *went*. At the same time, as the great early twentieth-century linguist Edward Sapir²⁸ aptly said, all grammars leak; all grammatical systems are

full of exceptions. Reconciling these two opposite tendencies is a big part of the acquisition process. The standard view on how this is done says that the child has a bias in favor of making generalizations (in the form of rules that correlate related groups of words with one another), as well as the ability to store irregular forms in long-term memory. As a child, you don't memorize *talked* (you only memorize *talk* and the rule "To form past tense, add *-(e)d*"). The more predictive value a rule has, the more "productive" it is said to be; the past-tense *-(e)d* rule is very productive, but not exceptionless: you must memorize *went*, *brought*, and *sang*, in addition to *go*, *bring*, and *sing*.

Here I would like to make a case for a stronger rule-bias, again based on arguments by Charles Yang.²⁹ The initial observation I want to build on is the fact that many irregular forms are actually quite regular. Consider the irregular past-tense forms of *bring*, *teach*, *seek*, *catch*, *think*, and *buy*. Although the verbs themselves have nothing in common, their irregular past-tense forms clearly do: *brought*, *taught*, *sought*, *caught*, *thought*, and *bought* can all be captured by an obvious rule. The same could be said of the past-tense forms of *sing*, *spring*, *ring*, and so on. As a second-language learner of English, I vividly remember trying to associate all the verbs to the mini-rules governing "irregular" past tense formation in the language. Along with Yang,³⁰ I suspect that this is also what children (unconsciously) do when they learn their first language. Instead of memorizing irregular forms one by one, as the standard model would have it, they most likely memorize which verb goes with which of the semi-productive mini-rules. In other words, we can think of the child learning irregular verbs as a rule-machine. Starting with the most general form (in the case of English past tense, "add *-(e)d*"), the child progresses until a form that does not fit the pattern is encountered. To be of any use, the rule-bias must be supplemented with a tolerance principle that prevents the child from abandoning the rule upon encountering the first form, or the first few forms, that don't fit. Let us tentatively assume that upon encountering these irregular forms, the child unconsciously builds representations (mental structures) like "for past tense, add *-(e)d*, except for the following verbs: *x*, *y*, *z*." Clearly, there must be a threshold beyond which the child realizes that the hypothesized rule is much more restricted than she thought. A rule that says "do *x* to all *y*'s, except *w*" is ok, but a rule like "do *x* to all *y*'s, except *a*, *b*, *c*, *d*, *e*, *f*, . . ." isn't. As the list of exceptions grows, the child is inclined to look for another pattern; specifically, a pattern among exceptions. So the child proceeds to hypothesize a second rule, less general than the first, but one that nonetheless captures many "exceptional cases." That rule, too, will come up against forms that don't fit, and the child will again be inclined to look for a pattern (a rule), a less productive rule, but a rule nonetheless, until, at last, the last rule hypothesized captures only one or two forms (which essentially amounts to saying that a handful of forms are stored as such). The model that emerges is one of a cascade of rules, each one of which is less productive than the last: "Do this to all *x*'s, except for *w*, *y*, *z*, where you should do *a*, except for *b* and *c*, where you should do *p*, except for *q*, where you should do *k*." We can think of this rule-bias (as Yang does) as a space-saving method,³¹ according to which the child

seeks to encompass as many forms as possible under a single rule/process, until the rule becomes too cumbersome (contains too many exceptions). But at that point, the child does not rush to memorize the exceptions one by one. Instead, she proceeds to consider a second rule that would reduce the memory load.

This way of thinking about how the child deals with irregularity and semi-productivity in language stresses the fact that the learner is biased toward rules, patterns, mental structures, and resists brute-force memorization as long as possible. I should add at this point that children seem exquisitely adept at keeping track of exceptions. It is true that many children overgeneralize the most general rule (such as “add *-(e)d*” for past tense), but they rarely if ever extend the use of a less general rule to an item that it doesn’t fit.³² Thus, forms like *goed* (instead of *went*) are frequent, but (contrary to what many might think) forms like *blought* (incorrect past tense of *blink*, on analogy to *think/thought*) are not, which again shows that children don’t simply learn by (unrestricted) analogy.

2.3 The need for cues

If a child just paid attention to English compounds, Noun–Noun combinations like *peanut butter*, *jelly sandwich*, *shoe box*, *paper towel*, etc., she might be led to think that English is in fact like Japanese in placing the most important member (the head) of nominal constituents last. Peanut butter is a kind of butter, a shoe box is a kind of box, a jelly sandwich is a kind of sandwich, and so on. This would clearly be the wrong move for the English child, because English constituents are overwhelmingly head-initial. The child’s attention must therefore be turned to more “revealing” examples, examples that would allow her to choose the right value for the parameters. To ensure this result, a few linguists have explored the possibility that each parameter comes with a special structural signature or *cue* – a special bit of data that would catch the learner’s attention and lead her to the right hypothesis.³³ If correct, this proposal entails that our biological endowment for language is richer than we thought. Not only do we come equipped with universal rules, and with parameters, but also with ways of knowing which bit of data to look for to set the values of these parameters. Needless to say, the cues should not be specific to the language being learned; they should be abstract templates. They must look like “If I encounter *x*, then my language is of this type, even if I encounter sentences that would suggest otherwise.” Thanks to cues, the learner essentially disregards (i.e., filters out) the misleading evidence, which is why all the cues proposed in the literature are unambiguous; they irrevocably favor one option (parameter value) over the competing hypothesis. Proponents of the cue-based approach³⁴ never fail to point out that in other cognitive domains, such as the visual system, neurons that are sensitive to highly specific features of the signal (such as horizontal lines, vertical lines, etc.) are known to exist;³⁵ so why should language not be assumed to be equipped with such sensitive mechanisms for data processing?

Proponents of cue-based learning have also pointed out that it is not enough for linguists to identify unambiguous data points that could serve as cues, such cues should also be accessible to the child, if the learning model is to be realistic. Linguists have access to all kinds of data, and can construct rarefied examples (experiments) that are plausibly not part of the input a normal child receives. The right cue should not be one that can only be found at the end of sentences that minimally contain, say, five clauses embedded one inside the other. Most, if not all, children would miss the cue, and thus, fail to learn. For this reason, David Lightfoot,³⁶ one of the strongest advocates of cue-based learning, has proposed a “Degree-0” restriction on cue-location, according to which cues must be located in unembedded contexts. That is, they must be found in the simplest sentences. In other words, they must be easily and quickly detectable.

Cue-based learning is without a doubt the safest parameter-setting method we have. But it is not the only model that has been explored. Janet D. Fodor³⁷ has argued that the effect of cue-based learning could be captured without assuming that cues are part of our innate endowment. Specifically, she has suggested a mechanism according to which the learner parses the input using all possible values of the relevant parameters, and rewards those values that are crucially needed for the sentence to be processed, and penalizes those values that would cause processing to halt. After enough examples have been encountered that crucially require one value, but not the other, the child ends up choosing this value as the correct one for the language of the environment. Fodor’s model mimics cue-based learning in that every time a sentence must be parsed that contains what cue-based proponents would call an unambiguous cue, one value of the relevant parameter will be rewarded and the other penalized in Fodor’s model, but this result is achieved without positing any cue in advance. Cues emerge as those data points that crucially need a specific value of a parameter to be processed. The burden of explanation is then on the parsing algorithm. For the model to be successful, it is crucial that the child’s parsing ability has matured enough so as to be able to reach the relevant portion of the sentence that would act as critical evidence. If the young child can only parse, say, four words, and the relevant bit of data with the cue-effect comes after the seventh word, the child won’t be able to set the parameter correctly until her parsing ability improves. In this sense, Fodor’s model expects the acquisition task to be slower than proponents of cue-based learning, for whom the learner need not be able to process sentences until she hits the cue. It is enough for the child to detect the cue, regardless of what else they understand in the sentence. The key, then, is to figure out which portions of sentences can act as cues, and to determine for each cue whether the child’s linguistic capacity is mature enough to reach these. The issue is thus one of biological maturation. (Note that Fodor’s approach predicts a certain parameter-setting path, where some parameters are set before others based not on some pre-determined order, or hierarchy of parameters, but on how the language faculty and the other cognitive systems supporting it develop.)³⁸

2.4 Frequency

Fodor's model, as presented here, also emphasizes the importance of frequency. It is only when enough data points of the right sort have been encountered that the choices are made. We could, of course, say that one relevant example suffices, but, given the noisiness of the data, I think it is safer to require a certain amount of evidence before choices are made. The gradual decline of the frequency of critical evidence, as opposed to the total disappearance of critical evidence, seems to me to be a more pervasive cause of language change (conceived of as development of a distinct parameter-setting configuration from the system internalized by the speakers in the environment – beginning with individual learners, but then spreading across the population as the percentage of learners of the new variety increases).

Looking at the frequency distribution of critical data not only can explain what choices learners make (and when), but it also reinforces the state of affairs Chomsky characterized as the poverty of stimulus (cf. Chapter 3). If what counts as critical evidence for a specific aspect of our linguistic knowledge is so infrequent, it won't do to say that that aspect of knowledge could have been learned from the input. If the evidence is too infrequent, it's as if it didn't exist, so the relevant bit of knowledge must be assumed to be part of our innate endowment.³⁹

Consider, for example, the case of subject-auxiliary inversion mentioned in Chapter 3.⁴⁰ One may entertain the possibility that the rule is one that is learned on the basis of examples like *Is the man who is tall happy?*, which show which auxiliary is picked in the case of inversion. Since the child will never encounter cases like *Is the man who tall is happy?*, one could imagine that the child picks the right option on the basis of exposure alone (of course, the child would still have to be endowed with the possibility of forming the right hypotheses, but I leave this aside for the issue at hand). But before we consider this possibility further, it is important to ask how often examples of the right sort (as complex and as telling as *Is the man who is tall happy?*) occur in the input that one would consider typical of a normal child learning English. A search through a database of child-directed speech has revealed that such examples are vanishingly rare, so rare that one can safely assume that many children proceed to develop the aspect of knowledge at issue without encountering them. The stimulus is indeed as poor as Chomsky claimed.

3 Navigating the Linguistic Space

So far the emphasis in this chapter has been on the learning biases that must be assumed, above and beyond the parameters that define the hypothesis space, for the child to ultimately acquire the language of the environment. But I would not like to leave the reader with the impression that once equipped with the biases discussed here (and no doubt many others) the child will automatically settle for

the right choices. Although the child typically ends up acquiring the target language, there is evidence that the learner is far more active – not only in the way she filters a vast amount of the input, but also in the way she explores possibilities that are not reflected in the environment. Not only does the child overgeneralize in a variety of circumstances, she also produces (for a while, at least) expressions that are simply not part of the language she is exposed to.⁴¹

Recall that even if we don't know exactly how many parameters Universal Grammar contains, we can safely assume that the hypothesis space the child operates with is fairly large. Finding the right combination of parameter values is non-trivial, and, as we may expect, the child makes mistakes. But the mistakes children make (if we ignore performance errors like false starts, etc. – which are just as common as in the speech of adults) are highly revealing. It turns out that the vast majority of errors that children make mimic the variation found across the world's languages, as the comparative method helps reveal. In other words, although children often fail to speak the local language during the course of language acquisition, we have no evidence that they fail to speak a human language at any point during the learning process. Put differently, children don't always map the input onto the right options, but even when they don't, they remain within the bounds of what is made possible by Universal Grammar. This claim is sometimes called the Continuity Hypothesis,⁴² which states that in order to account for the fact the children never step outside the bounds of what Universal Grammar makes available, it must be the case that the system that the child operates with is of the same nature as the one responsible for adult expressions in the world's languages.

Part of the excitement in studying language acquisition comes from identifying similarities between deviant expressions produced by the child and legitimate structures found across languages (but not in the target language). Among the many examples available in the literature, I would like to mention a few discovered by Rosalind Thornton, who has done seminal work in this area.⁴³ It is important to stress that the mistakes about to be mentioned have been frequently observed by researchers, though not every child makes them. In fact, no two children make exactly the same set of mistakes; the important thing is that whenever a mistake is made, it is possible to identify a language across the world in which that option is not in fact a mistake.

One of the mistakes English-acquiring children often make concerns subject-auxiliary inversion with questions introduced by the question word *why*.⁴⁴ Instead of the adult form *Why did you kiss her?*, children often go through a stage where they produce *Why you kissed her?* Crucially, they do not make the same mistake with other question words like *who*, *what*, *when*, etc. *Why* is special in this sense. The remarkable thing is that the same distinct behavior of *why* is found in adult languages. For example in Italian, the equivalent of “why,” *perché*, is the only question word that does not require subject-auxiliary inversion.⁴⁵ All other question words do. So when the English-acquiring child fails to invert the subject and the auxiliary in *why*-questions, she is testing an option that Universal Grammar makes available; she is, as it were, speaking a piece of Italian. So far as one can tell, there

is no evidence in the input to the child that would lead to this wrong choice. It's just a possibility that the child is freely exploring.

Likewise, many English-learning children repeat the question words in long-distance questions (so called because the question word must move a long distance, from within an embedded clause) like (adult form) *What did you say that Mary saw?*⁴⁶ Many children produce *What did you say what Mary saw?*, where the question word is repeated at the beginning of the embedded clause. Adult English rules out this option, but many other adult languages form long-distance questions this way: many dialects of German, for example, do exactly this.⁴⁷ So for a while many English-learning children speak a piece of German.

Examples of this sort of deviance abound. They are often easy to miss because most of the time the child produces forms that match the adult language of the environment, so it is easy to dismiss points of departure as just as few mistakes. But such mistakes are gems; they indicate that the child does not wait passively until she gets the right piece of critical evidence (or enough of it). The child does not mimic the adults around her; she actively explores the territory, constrained only by the system that makes this exploration possible in the first place: Universal Grammar, to which indeed all roads lead.

In closing this chapter I want to stress how hard the learning task (for language, and no doubt for many other cognitive domains) is, and how necessary it is for the child to be equipped appropriately. In order to navigate the hypothesis space provided by Universal Grammar, the child must be endowed with biases that channel her experience in highly specific ways. It is these biases that structure the input, and that result in a far more minimal (but far more informative) data intake than one may have guessed.

Learning from a massively ambiguous, noisy data set is bound to make the task one of approximation rather than perfect replication. When everything goes well, the child acquires a grammar that is sufficiently close to what is being used in the environment that communication can proceed, but sometimes (as when the ambient language is a pidgin), a new language emerges. Learning is a highly selective task;⁴⁸ the child can learn only what her biology renders recognizable (as Plato would have said, learning is remembering), and it also involves a lot of filtering – both filtering out wrong parameter settings from the innately available options, and filtering a massive amount of data, only a small fraction of which the child actively uses to make her choices. Even if learning undoubtedly involves domain-general processes, the mere fact that learning succeeds reinforces the need for highly specific mechanisms that both guide and rein these domain-general processes.

As learning takes place, the child becomes a better learner, as she narrows down the path of possibilities at every stage, and is able to use more mechanisms that reach maturation. Still, one should resist endorsing Piaget's view of development where later stages are constructed on what is acquired in earlier stages.⁴⁹ The right metaphor is not construction, but growth;⁵⁰ as the organism matures, more and more of its innately specified capacities become available.

As Elan Dresher points out,⁵¹ the acquisition task closely resembles a treasure hunt, where the child receives, from her biology, a map of the terrain to explore, a series of questions to answer, a list of answers to choose from, a series of steps according to which the search must proceed, well-defined clues to help her, and the possibility of exploring a few paths along the way, without wandering too far off track.

PART III

The Mental Foundations of Behavior

7

Making Sense of Meaning: An Instruction Manual

1 Meaning Inside the Head

Meaning is one of the most exciting issues of cognitive science; philosophers, psychologists, linguists, neuroscientists . . . everyone would like to understand how we make sense of the world around us, what thought is, and what sort of things we convey using language. Meaning has been called the “holy grail of cognitive science.”¹ Perhaps because so much is at stake, progress in this area has been maddeningly slow. The gulf between what we know and what we would like to know is immense.

As far back as Plato and Aristotle, philosophers have uncovered wonderful things we can do with language, what words can mean, and so on, but these nuggets of data have resisted deep explanation to this day. Consider the fact that when I say *John painted the house brown*² I understand this to mean in the usual case that John painted the exterior of the house brown, not its interior, but if I say *John painted the cave brown*, I understand this to mean that John painted the inside of the cave brown. Likewise, when a child hears the word *dog*, she automatically understands this word to apply to Fido and similar animals, not Fido alone, even if the word was first uttered in the context of Fido’s presence; and the child never entertains the mistaken notion that *dog* means ‘pet viewed from the front.’³ The child knows what to ignore. We spontaneously assign meanings to linguistic expressions in ways that are both highly constrained, and highly creative. How exactly this is done is far from clear.

It didn’t help that inquiry into meaning for much of the twentieth century (and still to this day) was heavily influenced by a logicist tradition that, much like behaviorism, looked down on psychological concerns, and banished talk of the mental.⁴ Influenced by Gottlob Frege and Alfred Tarski,⁵ semanticists have developed theories of linguistic meaning where words refer to entities in the world, and sentences are either true or false, based on what one finds outside the speaker’s mind. It didn’t matter that both Frege and Tarski warned against using the tools

of logic to study natural languages in their own works. Donald Davidson⁶ and Richard Montague,⁷ and many after them,⁸ made the bold claim that the apparatus to stipulate a semantics for a formal language (truth-values, entities in a domain, functions from such entities to truth-values) is appropriate to capture the semantics of natural languages. In being so focused on truth(-conditions) and reference, semanticists have allowed themselves to ignore some of the most fundamental insights of the early modern scientists like Descartes and Cudworth, who stressed the contribution of our mental organs in the way we see and think the world. Logical truths are mind-independent, but natural language semantics isn't.⁹

As a result of this logistic tradition, the line of inquiry that one could call "I-semantics"¹⁰ (with the stress on *Internalist* (mind-dependent) and *Individual* (creative), as opposed to *externalist* (mind-independent) and *social/communal* (normative)) has not yet amassed a comparable body of doctrine to that found in I-syntax, I-morphology, and I-phonology (the sort of findings discussed in Chapter 4). Nevertheless, a few aspects of linguistic meaning have recently been illuminated by approaches that acknowledge the centrality of mental structures, and stress the need to study "negative facts"¹¹ – ways in which we *can't* understand some expressions, not just the many ways in which we *can* understand them.¹²

Such themes are what I will amplify in this chapter, beginning with what meaning certainly isn't.

2 A Temptation to Resist at All Costs

At first blush, the following feels right. The word *cat* refers to the sort of animal that Fido hates, and the class of things that Felix is a member of. And I understand the sentence *Potatoes are in the closet* as true if indeed there are potatoes in the closet, and as false, otherwise. It is, thus, tempting to say that words like *cat* refer to certain entities in the world, and sentences like *Potatoes are in the closet* refer to states of affairs that are either true or false. Let us then say that that is what the meanings of words and sentences boil down to.

Doing this, however, would be confusing meaning and use.¹³ It is true that I can use the word *cat* to refer to Felix, and that I can say *Potatoes are in the closet* to indicate that this state of affairs indeed obtains, but we shouldn't conclude from this that meaning is reference, and that understanding sentences means knowing the conditions under which they may be true. One shouldn't confuse knowledge of language and knowledge about language; that is to say, one should distinguish between what speakers know by virtue of having a language faculty, and what they know by virtue of being a language-user in some environment. Knowing what the word *gold* means and knowing in what circumstances to use that word are very different things. (I may not know what the correct usage of *arthritis* is, but the word certainly contributes to how I understand *John suffers from arthritis*, so the word means something to me.)¹⁴

Making truth-evaluable assertions, referring to things, and inferring from statements and questions are some of the things we can do with words and sentences, but this is highly context-dependent, and thus highly variable. Already here we should be skeptical about taking reference and truth-conditions as the key notions for natural language meanings. Scientists (correctly) tend to shy away from variable facts of this sort, and instead they try to uncover stable phenomena; simple, tractable things they can hope to render intelligible. As Paul Pietroski¹⁵ stresses, whether (an utterance of) a sentence is true or false is what is known as an “interaction effect”; it’s the result of a host of factors, a hodgepodge of psychological states and physical circumstances that one may expect will forever resist interesting theorizing, the same way specific trajectories of apples falling from trees or stones thrown at a target will forever resist interesting theorizing. Physicists are right to abstract away from them; semanticists should do the same.

This would not bring the study of linguistic meaning to an end, for there is evidence that words and sentences have intrinsic semantic properties that are as stable as syntactic or phonological properties. Just like you know that *glbro* isn’t a word of English, and that *Who did Mary arrive after Bill kissed?* isn’t a good sentence of English, you also know that *John was eager to please* can’t mean (can’t be understood as) “John is eager for us to please him.”¹⁶ No context will change this fact. If words made direct reference to things in the world (what is often called “the universe of discourse”), and sentences had truth-conditions, semantics would depend on the varying character of communicative situations, but many facts about the way we understand sentences point to the fact that they don’t. Accordingly, we should separate as best we can questions about what language is used to do from questions about the means the language faculty provides for doing it; in other words, we should, as Chomsky has encouraged us over the years, draw a sharp dividing line between the functions and purposes of discourse and communication, and the mental devices that underlie our use of language to serve those functions and accomplish those purposes.

The ways in which we can assign meanings to expressions is a constant source of wonder. Consider the following sentence:¹⁷

(1) I almost had my wallet stolen.

You spontaneously interpret it to mean something like “I came close to suffering the theft of my wallet.” But upon reflection, the sentence could also mean “I was on the point of engineering the theft of my own wallet”; the same way we interpret *I almost had John arrested*. Admittedly, interpreting *I almost had my wallet stolen* in this way is not very natural, but this is likely due to the fact that it is somewhat bizarre to be interested in getting someone to steal one’s own wallet. The sentence at issue also has a third interpretation, which becomes clear when we relate it to *I almost had the game won*, which means something like “I was on the verge of being certain to win.” Accordingly our wallet sentence can also mean “I was on the point of being in a position where I would be certain of stealing

my own wallet.” This third interpretation is even more bizarre than the second: how can one come to stealing something that’s already one’s own property? But the interpretation exists, just awaiting a suitable context to be brought to light. (One only needs to think about fairy tales and sci-fi scenarios to realize how it is possible for us to conjure up meanings that the actual circumstances of our world don’t support easily.)

3 The Importance of Negative Facts

Because of the many ways in which words and sentences can be used, it is easy to forget the many ways in which they can’t be. While it is important to capture the fact that *The CEO called his secretary from Boston*¹⁸ can mean that either the secretary was from Boston, or the call was made from Boston, it is also very important to understand why the sentence can’t mean that the CEO was from Boston (even if it may be the case in the actual world).

Negative facts like the one just mentioned bear on theories of meaning, since the ways in which humans fail to associate signals with interpretations may well reveal important aspects of how humans understand language. In fact, negative facts may be the best data for understanding how we mentally structure linguistic signals and make sense of them. It’s such negative facts that also suggest that meaning is not just use, as Wittgenstein¹⁹ would have it.

Consider²⁰ the fact that if I say that Sue broke the vase at midnight, it necessarily entails that Sue broke the vase;²¹ and, for that matter, that the vase broke.²² But saying that the vase broke does not entail that someone broke the vase.

In a similar vein, if I say *every girl swam*, it means that every tall girl swam, which is different from what happens when I say *most girls swam*, from which it does not follow that most tall girls swam. And, if I say that every child ate, I necessarily mean that every entity that is a child is an entity that ate.²³

A theory of meaning should also capture the fact that when I say *Brutus stabbed Caesar*, I understand this to mean that Brutus was the “stabber” and Caesar the “stabbee,”²⁴ and not the other way round. Examples of this sort could be multiplied at will (think of the ways we understand, and don’t understand, *John persuaded Bill to leave* vs. *John promised Mary to leave*).²⁵

The importance of entailments of this sort is best illustrated in the context of little words like *any* and *ever* that linguists call “negative polarity items.”²⁶ Such items are only happy (or in linguists’ jargon, “licensed”), it seems, in the presence of negation. So, I can say:

- (2) a. Ginger won’t ever bite you.
- b. Nobody has any money.

But I can’t say:

- (3) a. *Ginger will ever bite you.
b. *John has any money.

Why is that? Part of the answer lies in the observation that negation is not the only thing that can license negative polarity items. Consider the following:

- (4) a. Does Ginger ever bite?
b. Does John have any money?
c. If Ginger ever bites Mary, I'll throw him out of the house.
d. If John has any money, he donates it to charity.
e. Every dog that has ever bitten Mary was named Ginger.
f. Every man that has any money is happily married.

Subtle changes, however, cause trouble for negative polarity items, as the bad sentences in (5) reveal:

- (5) a. *Some dog that has ever bitten Mary was named Ginger.
b. *Some man that has any money was lucky enough to marry Sue.

To figure out what is going on, let us start with “quantifiers” like *every* and *some*. Quantifiers are a bit like transitive verbs:²⁷ they require being associated with two types of elements, and take the following form: [[Quantifier *x*] *y*], where *x* is a noun and *y* a predicate, as in *Every student smokes*.

Some quantifiers allow a negative polarity item inside the *x* position, as in (6).

- (6) Every dog that has ever bitten Mary was called Ginger.
No dog that has ever bitten Mary was called Ginger.

Some quantifiers disallow this:

- (7) *Some dog that has ever bitten Mary was called Ginger.

Other quantifiers allow a negative polarity item inside the *y* position:

- (8) No dog has ever bitten Mary.

But many other quantifiers don't allow this:

- (9) *Every dog has ever bitten Mary.
*Some dog has ever bitten Mary.

The key to this puzzling state of affairs (which nonetheless poses no problem for children acquiring English) is to look at entailments. Saying *I have a dog* necessarily means that I have an animal (it entails a superproperty: *dog* is an instance of

“animal”) but does not necessarily mean that I have a poodle (it does not entail a subproperty: *poodle* being one kind of “dog”). Interestingly, words like *not* reverse entailments: If I say *I do not have a dog*, it necessarily means that I do not have a poodle, but it does not necessarily mean that I do not have an animal. For this reason, words like *not* are said to be downward entailing because they allow an inference from a property to a subproperty (the simple sentence *I have a dog* is upward entailing, allowing an inference from a property to a superproperty).

Going back to quantifiers, it turns out that some quantifiers display different entailing properties depending on whether you look at the *x* or the *y* elements they combine with. So, *every* is downward entailing with respect to *x*, but not with respect to *y*. When I say *Every dog barks* it necessarily follows that *Every poodle barks*. But saying *Every dog barks* does not necessarily mean that *Every dog barks loudly*. *Every* is in fact upward entailing when it comes to the *y*-element, as *Every dog barks* necessarily means that *Every dog makes noise*.

Turning to the quantifier *no*, the reader can check that it is downward entailing for both *x* and *y* (no dog barked \rightarrow no poodle barked; \rightarrow no dog barked loudly). By contrast, the quantifier *some* is upward entailing for both *x* and *y* (some dog barked means that some animal barked, but not that some poodle barked; it also means that some dog made noise, but not that some dog barked loudly).

I hope the reader sees the pattern emerging: negative polarity items are happy in downward-entailing contexts, not elsewhere. We would of course like to know why this is the key to their happiness (Paul Portner²⁸ notes some semanticists speculate that this has to do with the fact that negative polarity items express things like “the smallest amount,” which has a big effect on interpretation when the inference goes downward, but barely makes a difference when the inference goes upward), but the point of this discussion is that facts like negative polarity item environments crucially rely on properties that are intrinsic to sentence meaning (the directionality of entailment imposed by quantifiers), and do not resemble anything like the constantly shifting contexts of use.

4 Conceptual Instructions

Inevitably when semanticists study these facts independently of the contexts of use, the hypotheses that are proposed involve sentences being structured and how these structures support or impose certain modes of construal, but not others. Semantics thus becomes a species of mental syntax/chemistry,²⁹ the same way what we call “narrow” syntax and phonology are. Semantics becomes another aspect of the study of mental structures, and facts about meanings become reflections on the nature and operations of our language faculty, not about the way we interact with the outside world.

The best hypothesis we currently have about what the meanings of words and sentences are treats meanings as “instructions to build concepts,” to activate concepts we have independently of language (and that we may share with other species), or

build new ones that would be (literally) unthinkable without language.³⁰ Words and sentences can thus be thought of as procedures that impose certain mental traffic patterns among concepts; that is to say, they enforce perspectives on the way we think, and on the way we judge things as true or false (without determining truth or falsity), ways that are in large part as unique to us as language itself is. I will have much more to say about this hypothesis in the next chapter, when I examine the venerable theme of the relationship between language and thought, and show the sort of concepts that language makes possible. For now, I think it is more important to stress the need to understand the enterprise we call semantics in an internalist, psychologicist way. Instead of thinking of words as referring to external entities, the reader should think of them as commands to activate certain mental concepts and combine them according to instructions implicit in the structure of sentences.

Let me hasten to add that, of course, concepts ultimately connect to things in the world, but the way they do so is much more indirect than one might think, and so many concepts really don't seem to have any obvious referents out there in the world. I am not just thinking here of fictional characters like Hamlet or Ulysses, who don't exist in our world. I'm thinking of more mundane cases like "The average man,"³¹ or even words like *book*.³² Consider the latter. I have no problem saying *The book that Joyce wrote was too heavy for me to carry out of the library, but it doesn't matter because that book is unreadable anyway*. Does *book* refer to a physical object that Joyce produced, or just one of many copies in the library? And whichever it is, how can such an object be unreadable? Its content might be, but "it" can't, and yet *Joyce's book is unreadable* is perfectly intelligible. As we can see, it is very hard to pin down what simple words like *book* refer to (and I have said nothing about the book that's in my head and that I will likely never write). To make this point, Chomsky often uses the example of *London*, of which it can be said that it is too ugly, polluted, and angry, and should therefore be rebuilt elsewhere. Does *London* refer to a set of ugly buildings, or to the polluted air above it, or to its angry inhabitants, or to everything at once? And what does it mean for a place like London, distinct from another place like Cambridge, to be rebuilt some place else and still be (called) London and not something else?

The same effect can be observed with the run-of-the-mill adjective *green*,³³ which refers to very different things in *The house is green*, *The banana is green*, *The stoplight is green*, and *The country is green*.

As a final example of the interpretive richness made available to us, consider compounds involving the word *cookie*:³⁴

- Christmas cookie (cookie made to be consumed at Christmas)
- Yellow cookie (Yellow-colored cookie)
- Girl Scout cookie (cookie sold by Girl Scouts)
- Oatmeal cookie (cookie made of oatmeal)
- Walmart cookie (cookie sold at Walmart)
- Fortune cookie (cookie containing a fortune)
- Doggie cookie (cookie to be eaten by dogs)

We have no problems assigning stable interpretations to *book*, *London*, *green* and *cookie* even if their possible referents keep shifting, which strongly suggests that interpreting words is an internal affair; it's all inside the head. This is a conclusion that advocates of innate ideas like Descartes and Leibniz had already reached centuries ago. They knew that we understand words like "triangle," even if triangles as we understand them can't exist outside of our minds; we know what figures count as triangles, and we can tell them apart from squares and circles. We know what their properties are even if we also know that no "real" object has these properties. What the word "triangle" does is illuminate this concept the way a flashlight can reveal a painting otherwise trapped in the dark.³⁵

5 Cartesian Insights Once Again

Internalism about meaning has a distinguished tradition (obscured by the anti-mentalism of the twentieth century). As Chomsky³⁶ and McGilvray³⁷ have pointed out, the seventeenth-century philosopher Ralph Cudworth,³⁸ who we have already encountered in Chapter 2 in the context of *Cartesian Linguistics*, spoke of our "innate cognoscitive power" to form lexical concepts, and Humboldt appealed to our productive "mental instrument" to engender concepts.³⁹ Cudworth agreed that circumstances may occasion or invite a concept, and Humboldt said that a signal from a person's speech may cause a matching (but not completely identical) concept to be constructed in the mind of the hearer. But as Chomsky pointed out in his review of Skinner's book,⁴⁰ however incited and inclined in specific circumstances, a human being can always choose to say something different from what the listener may expect (or say nothing at all). No one really knows what the relevant trigger of a specific word use is, and perhaps no one ever will know (as Descartes suspected). Recognizing our mental machinery as essential is a very important, indeed crucial, step towards understanding, but how that machinery is put to use remains mysterious. Words are apt for use; they "guide" use⁴¹ – as Descartes might have said (according to McGilvray),⁴² our will's freedom to judge is constrained by some ideas of the understanding – but ultimately, how we exercise this capacity is beyond our understanding. That's our human nature.

In the words of McGilvray,⁴³

Our cognitive system, with language at its heart, is not only flexible enough to deal with a variety of circumstances, it can also detach itself from circumstances, and range widely. We can speculate, engage in wonder, idealize, construct thought experiments, and cultivate what Kant called a free play between the imagination and the understanding.

The human mind, empowered with this remarkable linguistic capacity, is designed to be creative, and is so even at the earliest stages of development (think about

how quickly children engage in fantasy, play, story telling, personification, etc.).⁴⁴ We are equipped with innate concepts that thanks to language can be combined in an endless fashion, and we make free use of this system.

This is what Chomsky called the creative aspect of language use.⁴⁵ By creativity, Chomsky does not necessarily mean poetry and science, Shakespeare and Newton, Goethe and Einstein; he means the sort of ordinary creativity we all engage in. As the Cartesians had already emphasized, ordinary language use is in Chomsky's summary⁴⁶ "typically innovative, guided but not determined by internal state and external conditions, appropriate to circumstances but uncaused, eliciting thoughts that the hearer might have expressed the same way." It's what makes ordinary language use unbounded, stimulus-free and coherent/appropriate.

The best one can do is characterize the machinery, the mechanism that makes this behavior possible, but it is unlikely there will ever be a science explaining this behavior.⁴⁷ Our creativity makes it possible for us to do science (Chomsky sometimes speaks of a science-forming faculty),⁴⁸ but our science may never be able to capture the creativity that enables it. Almost by definition, the detachment that makes it possible for us to freely exercise our language capacity is impossible to rein in via theoretical understanding.

Coming back to "meaning," one should bear in mind that a theory of linguistic meaning can't, and shouldn't, be a theory of everything that makes a sentence meaningful. Such "theories of everything" never materialize. I have put emphasis in this chapter on what sort of facts (especially what Pietroski has called the "negative facts") one should focus on to reveal the interpretive mechanisms humans make spontaneous use of, because even if it is true that specifying truth conditions was not always the only semantic project in town, this way of conceiving of the semantic enterprise has been so dominant in the last century that it is important to redirect inquiry toward the more mentalist concerns of the Cartesians.

Meanings are not normative, contingent, arbitrary, conventional, learned, or dependent on the outside world. They are like an instruction manual, a blueprint for conceptual assembly; like genes, which guide the building of the organism that interacts with the environment, meanings guide the construction of concepts that are put to use. This way of thinking about meanings is not only more in line with the project of revealing mental structures at the heart of cognitive science, it also forces us to study the systems with which the language faculty interacts. If meanings provide instructions to cognitive systems to construct (activate, combine, . . .) concepts, it becomes very important to know what those cognitive systems are (what the architecture of the mind is), and what they can do independently of language. This is what is called the study of the interfaces, and it is bound to reveal the specific contribution language makes to how we think.

8

Wonderful Mental Life: Unthinkable without Language

1 What Does Language Contribute to Thought?

Does a mind without language look more or less like one with language? Is the difference just a matter of externalization (pronunciation/signing)? If not, how are we to characterize the difference?

Most people interested in addressing such questions scientifically would, I think, readily agree that the best way to find out would be to examine creatures without language, such as our closest relatives (monkeys, apes), or even infants whose language faculties haven't matured yet (if we can ever get to such a stage, given how early infants display linguistic abilities; cf. Chapter 3), and probe what they think, and, more importantly, how they think it.

Long gone is the time when philosophers and psychologists believed that thought was the exclusive property of adult humans¹ (although this view is often attributed to Descartes, it is not at all clear that he held it;² he merely wanted to emphasize the gulf between us and other species). We now have overwhelming evidence that the mental life of non-linguistic creatures is extremely rich – distinct from ours, to be sure, but equally fascinating.³ In fact, when we compare some of the capacities that enter into acts of perception and cognition in humans and in other primates, we discover extensive similarities. Our visual and auditory capacities are not all that different, nor are the neural mechanisms underlying such capacities, as far as we can tell.⁴ As Darwin guessed,⁵ the psychology of humans is continuous with that of nonhuman animals, and we can be pretty confident that we have a lot of cognitive mechanisms in common.

But such similarities should not obscure the fact that our cognitive achievements differ strikingly from those of our closest relatives. As Elizabeth Spelke points out,⁶ although all animals have to locate and categorize food, we alone have developed the subtle chemistry (art and science) we call cooking. Although many juvenile

animals engage in playful behavior, only we came up with the idea of turning play into complex games. Likewise, although all animals must somehow understand critical aspects of the world they live in, only humans have gone on to do science, an activity that is concerned with things too small or too large to perceive. And, finally, despite the fact that many animals have a rich social life, we are the only species who codifies it into laws.

I could go on to discuss music, theater, architecture, mathematics, politics, religion, industrial revolutions, and many more outputs of our one-of-a-kind cognition. Herein lies the paradox: when we do comparative cognition, we are faced with striking cases of continuity and remarkable cases of discontinuity. The hypothesis I will explore in this chapter – based on research by developmental psychologists, philosophers, linguists, and anthropologists, and especially by Spelke⁷ – is that language is key to understanding this paradoxical situation. In other words, we may not need to posit many different mechanisms to account for human uniqueness. A single algorithm may be powerful enough to do the trick. That algorithm lies at the heart of our language faculty, and, as some have suggested, may be crucial to understanding the relation between language and thought. Let me stress right away that what follows is but one way of making sense of the result of comparative cognitive studies. Though attractive, the hypothesis that is (I think) slowly becoming the consensus may turn out to be wrong, but even if this turns out to be the case, discussing it in this book not only allows me to show how to resolve the paradox I have introduced, it also allows me to review some of the most robust findings in cognitive science (especially developmental and comparative psychology), and introduce concepts like modularity, which I think everyone would agree must somehow be part of how we think about the architecture of the mind. It will also help me refine some of the observations made about natural language semantics in the previous chapter.

2 Core Knowledge Systems

Let me begin with what we share, cognitively speaking, with other animals. Recall from previous chapters (especially Chapter 2) that just the very fact that we experience certain things and not others requires the presence of special-purpose mechanisms that perform specific tasks.⁸ Such mechanisms have been called mental modules⁹ or mental organs.¹⁰ We can also call them special knowledge systems. Nonhuman animals must be endowed with such systems too; they wouldn't be able to make sense of the world around them and survive without them. Because such systems appear to be fundamental to our cognitive activity, let us call them, with Spelke, "core knowledge" systems. Comparative and developmental research (i.e., research focusing on nonhuman animals and young infants) has already revealed a few such core knowledge systems, and their basic properties.¹¹ We still don't know exactly how many there are, but (Spelke claims,¹² and I concur) we can be quite confident that there aren't as many as proponents of so-called massive modularity¹³

claim (perhaps a hundred of them). We have very robust evidence for four or five core knowledge systems (there may be a few more, though):¹⁴ one system specializing in objects and their mechanical interactions (a system sometimes called “object mechanics”), another specializing in agents (animate things) and their goal-directed actions, a third concerned with sets and numbers (our “number sense”), a fourth dealing with places and geometric relationships (natural geometry), and more recent evidence suggests that there may be a fifth core knowledge system dealing with social partners, groups, and relations, and the way we understand other minds (theory of mind).

Such systems, to be described briefly below, are at the root of our capacity to form rudimentary theories of the world around us, what is sometimes called “folk science”; these theories are the foundations of physics (object mechanics; “naïve physics”), mathematics (number sense), biology (animate vs. inanimate beings; “folk biology”), navigation (natural geometry), and psychology/social science (theory of mind; “folk psychology”). These core knowledge systems give us an intuitive grasp of what is going on in each of these domains, but humans uniquely develop to transcend what Spelke calls the “signature limits” of core knowledge systems, and develop scientific, i.e., far less intuitive accounts of each of these domains. Core knowledge, however, gives us a head start, and enables us to measure the highs (and lows) of our cognitive achievements.

2.1 Thinking about objects

Consider first the object mechanics system.¹⁵ There is now massive evidence that infants come equipped with a system that enables them to perceive objects and their motions, to fill in the surfaces and boundaries of an object that is partly occluded, and to represent the continued existence of an object that moves fully out of view. Such a system accounts for hundreds of experimental results performed (and replicated) over the past 20 years. For example, such a system accounts for the fact that if a 5-month-old infant is presented with the following sequence of events – (1) a puppet on a stage, (2) a screen concealing that puppet, (3) a second puppet shown to the infant then moved behind the screen concealing the first puppet – the child will look longer (suggesting a violation of her expectation, hence revealing her belief) if there is only one puppet left on stage when the screen is removed. From this it is safe to conclude that children don’t think objects cease to exist when hidden.

On the basis of experiments of this sort, Spelke and her collaborators have proposed that human infants represent objects in accord with three spatiotemporal constraints on object motion:¹⁶

- 1 *Cohesion*: Infants represent objects as cohesive bodies that maintain both their connectedness and their boundaries as they move (two objects are not expected to fuse if they start as distinct objects; nor is an object expected to split into two as it moves).

- 2 *Continuity*: Infants represent objects as continuous bodies that move only on connected, unobstructed paths.
- 3 *Contact*: Infants represent objects as bodies that interact if and only if they come into contact.

Experiments with other animals (both adults and juveniles), such as adult monkeys or newly hatched chicks, reveal a similar pattern.¹⁷ Using similar experimental methods (such as preferential looking to detect violations of expectations), adult monkeys performed as well as, and sometimes exceeded, human infants. More imaginative methods had to be used in the case of newly hatched chicks, but the results confirmed that humans are certainly not the only creatures to represent objects as spatiotemporally continuous objects (sometimes called “Spelke-objects”). In other words, we share with many other species at least some aspects of how we reason about the physical world, which may not be too surprising since after all object perception is fundamental for survival. Thanks to this system, human infants as well as other animals meet the world with a series of fundamental expectations that effectively construct their own physical environments.

2.2 Thinking about numbers

In addition to being able to represent how objects interact, humans and other animals have the ability to represent several objects at a time, but only a few of them (three or four);¹⁸ thus, humans are able to attend to three or four separately moving objects, when the objects’ boundaries accord with the cohesion and contiguity constraints just discussed. Human infants, as well as nonhumans, fail to track objects beyond this set size limit, giving rise to surprising experimental results,¹⁹ such as the inability of young infants to prefer four vs. five cookies, or even one vs. five cookies. They can certainly prefer one vs. two, or two vs. three, or one vs. three cookies because they can keep track of how many cookies there are, but beyond three (or four), young infants perform at chance – as surprising as it seems to us.²⁰

The ability to grasp three or four objects at a time, without counting, just by perceiving and categorizing, is called subitizing,²¹ and it is often discussed in connection with our number sense, because it is one way of representing (small) numerosities.²² But we have a different system capable of dealing with larger numbers.²³

Numerous experiments have established the existence of a system that provides us with a sense of approximate numerical values and relationships. The performance of this system is characterized by Weber’s law:²⁴ as numerosity increases, the variance in subjects’ representations of numerosity increases proportionately, and therefore discriminability between distinct numerosities depends on their difference ratio; the closer the two numerosities are, the harder it is to tease them apart (this is the so-called distance effect); and the bigger the two numerosities are (keeping the distance between them constant), the harder it is to tease them apart as well (this is the so-called size effect).

Young infants (as early as 6 months old) are better at discriminating visual arrays of 8 dots vs. 16 dots than they are at discriminating between 8 dots vs. 12 dots, or 26 dots vs. 32 dots (infants are expected to look longer at the array with the novel numerosity in such cases).²⁵ Such findings have been replicated with tones instead of dots on a screen, and suggest that numerosity representations are not limited to a particular sensory modality (visual or auditory) or format (spatial vs. temporal). In other words, number representations are abstract. They have also been shown to be manipulable by addition and subtraction.

A similar (imprecise) capacity to represent numerosities has been observed in nearly every animal tested, from fish to pigeons, rats to primates, pointing to a near-universal number sense.²⁶

2.3 Thinking about actions

A third core system represents agents and their actions. The spatiotemporal principles we reviewed above in the context of objects do not govern the representations of agents – entities that need not be cohesive, continuous in their paths of motion, or subject to contact in their interactions with other agents. As Kinzler and Spelke²⁷ review, experiments have revealed that infants expect that agents have intentional actions directed toward goals (infants do not interpret the motions of inanimate objects as goal-directed), that their goals are achieved through means that are efficient, and that agents interact contingently and reciprocally. Furthermore, although agents need not have perceptible faces, when they do, infants (even newborns) use the direction of gaze of these faces to interpret their social and non-social actions.

Goal-directedness, efficiency, contingency, reciprocity, and gaze directions thus define agent representations, and not surprisingly nonhuman animals are sensitive to these properties²⁸ (they are sensitive to what predators can and cannot see, for example), suggesting that we are dealing with another dedicated system that is evolutionarily ancient and that persists over human development.

2.4 Thinking about space

The fourth core system I want to discuss is natural geometry. It would come as no surprise to Plato, Descartes, and Leibniz to see developmental psychologists arguing for this innate capacity. As I already pointed out in previous chapters, these philosophers never failed to use examples from the domain of geometry to motivate innate ideas. Children and nonhuman animals have taken part in numerous experiments that suggest from the way they navigate that their behavior is governed by Euclidean principles. Perhaps the most dramatic evidence for natural geometry comes from the literature on dead reckoning (aka path integration).²⁹ Desert ants, for example, leave their nests in search of food, traveling a long and tortuous path from the nest until food is found. At that point the ant makes a straight-line path

for home. Crucially, the path followed on the return differs dramatically from the outbound journey they took to find food, and does not appear to be guided by landmarks (if all potential landmarks are removed, the straight-line path of the ant remains highly accurate). It looks like the path is determined solely by the geometric relationships between the nest location and the distance and direction traveled during each step of the outgoing trip. It is safe to conclude from this that ants (and virtually all the other animals we know of) have a natural sense of geometry, which dwarfs my poor sense of direction.

2.5 Thinking about society

Finally, evidence is beginning to accumulate in favor of a fifth core knowledge system for identifying and reasoning about potential social partners and group members.³⁰ Infants show a visual preference for members of their own race; infants also look preferentially at faces of the same gender as their primary caregiver, and listen preferentially to speakers of their native languages. As Kinzler and Spelke suggest,³¹ these elements may form the basis for cultural learning, though more work in this area remains to be done.

What is clear is that some core knowledge systems guide the way humans and nonhumans interact, perceive, categorize, represent, and interpret the environment. (How else could any of us survive?) It is easy to see how all these systems may provide the basis for all the things that humans do, but, due to their signature limits, it is also clear that such systems fail to account for what we do. The clearest example of this failure, in my opinion, comes from mathematics. The core knowledge systems provide two ways to represent numerosities: accurate representations of small numerosities via subitizing, and less accurate representations of (large) numbers via approximation. What is missing is our (unique) ability to represent large numerosities precisely. Core systems can't achieve this, so what, in Spelke's terms,³² could make us so smart (or, in more neutral terms, so different)?

3 Modularity

Before addressing this question, I would like to list the characteristics of core knowledge systems that Spelke ascribes to them in her works.³³ Such systems are:

- domain specific: each serves to represent (categorize) only a subset of the entities in the child's (/animal's) surroundings
- task specific: the representations (/categories) constructed (made available) by each system guide only a subset of the actions and cognitive processes in the child's/animal's repertoire

- relatively encapsulated: the internal workings of each system are, as it were, inaccessible to other representations and processes found in other cognitive systems
- isolated: the representations that are constructed by distinct systems do not readily combine together

These four characteristics are among the telltale signs of what Jerry Fodor called mental modules in his influential book *The Modularity of Mind*.³⁴ Because the architecture of the mind defended by Fodor has been at the center of many discussions in cognitive science, I would like to take this opportunity to sketch Fodor's argument here, before returning to the relation between language and thought.

Fodor's central hypothesis is that there exist in the mind/brain mental faculties that are domain specific, genetically determined, likely associated with distinct neurological structures, and computationally autonomous. These Fodor calls modules. Crucially, this does not exhaust what the mind/brain is; there are likely to be processes that are not domain specific (e.g., a kind of central processing unit drawing information from the various modules),³⁵ but modules play a prominent role.

In some sense,³⁶ this view of the mind goes back to the nineteenth-century movement called phrenology and its founder Franz Joseph Gall, who claimed that the individual mental faculties could be associated precisely, in a sort of one-to-one correspondence, with specific physical areas of the brain. Hence, someone's level of intelligence, for example, could be literally "read off" the size of a particular bump on his skull. This simplistic view of modularity has been abandoned over the course of the last century, but some aspects of it are alive and well in modern cognitive (neuro)science³⁷ (I return to this in Part IV, especially Chapters 10 and 12). Fodor revived Gall's vision, but distanced himself from the notion of precise physical localizability of the mental faculties (without denying that they are somehow instantiated in brain matter), and instead used arguments from modern cognitive science to support his hypothesis.

According to Fodor, a module has some of the properties that behaviorists would have attributed to "reflexes." The key difference boils down to information. For the behaviorists no information was added when reflexes processed aspects of the environment. Fodor rejects this claim, as he takes modules to modify, transform, and interpret the signals they receive (think of the work that must be done to process a 3D object from a 2D signal on the retina). In other words, some form of intrinsic computation must be assumed for the modules.³⁸ But Fodor insisted that the modules are, like reflexes, cognitively impenetrable; that is, they are not influenced by other cognitive domains. For example, no matter how often you encounter the Müller-Lyer visual illusion discussed in Chapter 2, and reproduced here, your mind will be tricked into believing that one horizontal line is smaller than the other, even if you know that this is an illusion.



You can't help but see it exactly as you saw it the first time. This is taken by Fodor to indicate that other domains, including one's beliefs, cannot influence such reflex-like information-processing mechanisms that define modules.

Fodor lists the following properties for his modules:³⁹

- Domain specificity (modules only operate on certain kinds of inputs – they are specialized)
- Informational encapsulation (modules need not refer to other psychological systems in order to operate)
- Obligatoriness (modules process in a mandatory manner)
- Fast speed, probably due to the fact that they are encapsulated (thereby needing only to consult a restricted range of options) and mandatory (time need not be wasted in determining whether or not to process incoming input)
- “Shallow” outputs: the output fed by a module to more domain-general systems is very simple
- Characteristic mode of development (fixed developmental schedule)
- Distinct neural architecture (as suggested by highly specific pathologies)

Pylyshyn⁴⁰ has argued that while most of these properties tend to be found in all the modules suggested in the literature, one stands out as being the real signature of a module: the encapsulation of the processes inside the module from both cognitive influence and cognitive access. This is referred to as the “cognitive impenetrability” of the module. To illustrate this, consider the following anecdote:⁴¹ Ask someone to look at the clock on the wall and tell you what time it is; after they do so, ask them to tell you the shape of the hands of the clock, without looking at the clock again. You'll find out that people tend to get the answer wrong. They must have seen the hands of the clock (they processed that information) to tell you what time it was, but somehow that information is lost; it's inaccessible, encapsulated.

More recently, proponents of the trend in cognitive science called evolutionary psychology⁴² have claimed that modules are units of mental processing that evolved in response to selectional pressures in the course of history. On this view, much modern human psychological activity is rooted in adaptations that occurred earlier in human evolution, when natural selection was forming the modern human species. I will not be concerned with this particular view of modules. All I want to point out is that many characteristics of Fodor's modules can be found in core knowledge systems, which can be said to form some of the most basic organs of the mind.

This being said, core knowledge systems, while they reinforce the architectural claim made by Fodor in *The Modularity of Mind*, also point to the degree to which our human mind is not modular, since, as I pointed out already, building on work by Elizabeth Spelke and others, humans transcend the limits of core knowledge systems.

4 How Language Transforms Thought

The hypothesis I would like to discuss here is that natural language plays a significant role in breaking the bounds of modularity. Specifically, following a growing number of researchers,⁴³ I would like to suggest that language is able to combine information from various modules (core knowledge systems) that would otherwise be hopelessly isolated. This allows the formation of novel representations, novel molecules of thought as it were, that go (it seems to me) a long way toward explaining our actions and achievements. As we will see,⁴⁴ it's as if human language provides a mold or universal currency with which mental transactions across modular boundaries can be carried out without difficulty; thanks to language, apples and oranges can be mixed to form delightful cognitive cocktails.

Let me illustrate the cognitive influence of language on core knowledge systems by looking at results from the literature on spatial representation and on number, where the evidence is, I think, very strong.⁴⁵

4.1 Navigation

As we already saw, animals are endowed with a natural geometry module that enables them to represent space, and navigate through it. In so doing animals display remarkable abilities (cf. dead reckoning), but also quite revealing limits. For example, various animals can readily locate food by searching in a particular geocentric position (say, the northern corner of the experimental setting), or by searching near a particular landmark (say, right next to a box), but they have a lot more difficulty searching in a particular geocentric position oriented with respect to a particular landmark (say, north of the box).⁴⁶ So although animals can represent "north of the room" and "next to the box," they do not seem to be able to combine these representations. Other experiments suggest that animals often fail to call on representations that we have independent reason to believe that they have mastered when they are placed in a new context.⁴⁷

Spelke and colleagues⁴⁸ have used the same experimental setting for landmarks and geocentric positions with young infants, and discovered that young infants show the same inability to combine representations in novel ways. Infants, like animals, can learn to search directly at a geometrically defined landmark, and they can learn to search directly at a non-geometrically defined landmark, but combining these two sources of information into one poses problems. More interestingly, research⁴⁹ suggests that the transition to a more flexible navigation system is closely related to the emergence of spatial language, when children show clear mastery of spatial expressions involving terms like *left* and *right*. This shows a correlation between spatial language and flexible navigation, but it does not yet establish a causal link.

The way such a causal link is often suggested in the literature is via experiments that would interfere with the use of the source of the cause during a given task.⁵⁰ In this case, psychologists have tried to “block” the use of language during navigation experiments by means of a “shadowing” task, a task that is performed by the subject at the same time as the one at the heart of the experiment and that taxes attention and memory. They found out that if the shadowing task is verbal in nature, performance in the navigation experiment decreases for adults, who otherwise outperform young children and animals. But if the shadowing task is non-verbal, the flexible pattern of navigation attested in adults is unaffected. This suggests that language is the source of the flexibility that is not present in young infants and animals. Language appears to be the source of novel representations that enable linguistic creatures to think differently, more flexibly, and more creatively.

4.2 Numerosity

A similar influence of language on cognition can be found in the context of number. Recall from our discussion of the number sense that humans and other animals come equipped with two modes of numerical representation: a system that subitizes (represents three to four objects at a time; enabling children to recognize one object as distinct from another, and add or subtract one object from a small set), and a system that approximates (and deals with larger quantities). These two modes of representation appear to be distinct, and, experiments suggest, do not spontaneously combine in animals’ or young infants’ minds. The difference between the two systems at hand is at least threefold:⁵¹ (1) the subitizing system has a size restriction that the approximate system lacks; (2) the approximate system is subject to Weber’s law, while the subitizing system isn’t: discriminating between one vs. two objects is as easy as discriminating between one vs. three objects, even if two is closer to one than three is; i.e., there is no distance effect; and (3) representations of numerically distinct objects are robust over occlusion (even when hidden during an experiment, the child knows exactly how many objects there are), while approximate numerosities are not.⁵²

In sum, one system represents small numbers of persisting, numerically distinct individuals exactly, and takes account of the operation of adding or removing one individual from the scene, but evidence suggests that this system fails to represent individuals as a set; rather, it instantiates: one, another one, and another one, period. The second system represents large numbers as sets; it allows comparisons across sets, but it fails to represent sets exactly. In other words, as Spelke has argued,⁵³ infants and other animals represent “individuals” and “sets,” but seem to be unable to spontaneously form representations of “sets of individuals.”

The concept “set of individuals” is, of course, central to counting, elementary arithmetic, and our notion of natural number. We therefore expect young infants to have problems with natural number terms like “two.” Young infants should

also miss the point of the verbal counting routine (even if they mimic it), and indeed, a large body of experiments shows that young infants suffer from these two problems.⁵⁴

Experiments have shown that children's understanding of the counting routine develops in four stages.⁵⁵ At stage 1 (roughly, 2–2.5 years), children understand that “one” refers to “an object.” They also understand that all other number words refer to arrays with more than one object, but for them, two toys or six toys amount to the same thing: “more than one toy.” If asked to point to pictures with “two toys,” as opposed to pictures with “four toys,” they point at random. At stage 2 (2.5–3.25 years), children come to understand the meaning of two, as a set of one individual and another individual. Three and above refer to “sets other than one or two.” Three months suffice to reach stage 3, where children now grasp the meaning of three, as a set composed of one individual, another individual, and another individual. Finally, in stage 4, children come to understand that each number word designates “a set of individuals,” and that the set of individuals designated by each number word contains one more individual than the set designated by the previous word in the counting routine. To reach that stage, children must go through more than a year of development, during which they grow the means to combine the two distinct modes of representation of numerosity that they possess at birth. The key developmental stage is, of course, the step from stage 3 to stage 4. Until then children can rely on the subitizing mode to handle one, two, and three individuals, but once they are beyond the signature limits of the subitizing system, the inaccuracy of the approximate system is revealed. I think Elizabeth Spelke is right in assuming that language serves as the medium to reach stage 4.

Several experiments show that language is deeply involved in counting (try to perform even elementary arithmetic in a language other than your own native tongue, with large numbers, like adding 25 to 33). For example,⁵⁶ you can take bilinguals and perform experiments where part of the story (containing facts about number) is given to them in one language, and another part of the story (containing other facts about number) in the other language. Once subjects have shown that they have memorized each set of facts correctly, if you test them on all the facts, in both their languages, you will see that subjects retrieve facts about number from either language equally well, so long as numbers are either small and precise or large and approximate. But they show a marked preference for reporting large exact quantities in the same language as the one used in the original story. This suggests that approximate numerosities and small numbers are stored independently of language. But large exact quantifiers are tied to language, as one would expect, since large exact quantities require the subject to transcend the limits of the core knowledge systems, and rely on language to combine representations of distinct systems.

To conclude this section on the number sense, adult humans show evidence for three systems, or modes of representation, with the key addition being dependent on the language faculty.

5 What Makes Us Human

Aside from these results from the realms of navigation and number, there is also developmental evidence that children master the ability to reason about what other people think when they show evidence for the ability to embed sentences inside sentences containing verbs like *think* or *believe*.⁵⁷ Until then if you ask them *What did Mary think was in the box?*, they seem to understand this question as *What was in the box?* This would suggest that the rudimentary theory of mind, the core knowledge system that allows animals to reason about other beings, is also dramatically affected by core properties of the language faculty.

Although nothing in science should be taken for granted, there seems to be growing evidence for the claim that our distinct kind of intelligence (“thought”) is due to our capacity to flexibly combine representations that would otherwise remain isolated. Human language offers the clearest example of an algebraic system that combines (indefinitely). The mode of combination happens not to depend on its conceptual roots: we combine events, colors, numbers, emotions, and numerous other entities in sentences, and treat them all like x ’s and y ’s, thereby transcending the limits of domain-specific systems that can only combine an x with another x .

With language in humans, biology has found a way to develop a system that instructs the mind to fetch a concept here and fetch a concept there and combine them even if the concepts by themselves wouldn’t fit naturally; their word-clothings make them fit together. The roots of our knowledge are ancient, and continuous with other species, but our kind of thought, our creative bent, as it were, required the evolution of lexicalization, which applies a uniform format to concepts that would never combine otherwise. Human language, under this hypothesis, takes the form of a central processing unit that creates a *lingua franca*, a genuine language of thought, out of the mutually unintelligible dialects of thoughts that are the core knowledge systems.

Another metaphor for the cognitive effect of human language would be the Swiss Army knife. Until language emerged, the minds of our ancestors were full of various tools, each tailored to specific needs. With language, all these tools were combined into a flexible all-in-one tool that makes available a variety of solutions (tools) whose effects can be combined spontaneously.

Anthropologists⁵⁸ have noted that many aspects of religious beliefs consist in transposing an aspect of (cognitive) experience onto another cognitive domain, yielding novel combinations, creating novel inferences. They have also noted⁵⁹ that tools and cultural artifacts specific to *Homo sapiens* often involved the juxtaposition of familiar parts into an unfamiliar (novel) whole; for example, a sculpture fusing a human face, animal ears, and wings, which is meant to convey the thought that the resulting product is more than the sum of its parts.

I conclude with Marc Hauser’s characterization of what makes us unique (what he dubs “humaniqueness”) is the ability to:⁶⁰

- 1 combine and recombine different types of information and knowledge in order to gain new understanding
- 2 apply the same rule or solution to one problem to a different and new situation
- 3 create and easily understand symbolic representations of computation and sensory input
- 4 detach modes of thought from raw sensory and perceptual input

I have suggested elsewhere⁶¹ that these four distinct features of human thought boil down to our ability to lexicalize concepts and combine them freely. The creative aspect of language use so central to Descartes' philosophy has left footprints all over our cognitive world. As Hauser put it, while other animals have laser-beam intelligence,⁶² we are uniquely endowed with floodlight intelligence, which this chapter has suggested is due to our unique biological endowment: the language organ.

9

Grammar Caught in the Act

1 Competence and Performance

George Miller¹ once characterized the shift from behaviorism to modern cognitive science as the realization that behavior is simply the evidence, not the subject matter of psychology. To make this clear – that is, to stress that linguistic theory is concerned about the representations and computations that enter into language, not about sentences (utterances, behaviors of speakers) per se – Chomsky² introduced the famous distinction between competence and performance in the first chapter of *Aspects of the Theory of Syntax*, which in my view remains the clearest statement of the commitments of linguistics seen as part of cognitive science. As he puts it:³

Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance.

... To study actual linguistic performance, we must consider the interaction of a variety of factors, of which the underlying competence of the speaker-hearer is only one.

... We thus make a fundamental distinction between *competence* (the speaker-hearer's knowledge of his language) and *performance* (the actual use of language in concrete situations).

By this distinction Chomsky meant several things. First, like any good scientist, he wanted to approach a complex phenomenon like language by making a series of idealizations. The competence–performance distinction is one of the most basic (and necessary) idealizations one can make in this domain. As Chomsky notes:⁴

Only under the idealization set forth [above] is performance a direct reflection of competence. In actual fact, it obviously could not directly reflect competence. A record of natural speech will show numerous false starts, deviations from rules, changes of plan in mid-course, and so on. The problem for the linguist, as well as for the child learning the language, is to determine from the data of performance the underlying competence. Hence, in the technical sense, linguistic theory is mentalistic, since it is concerned with discovering a mental reality underlying actual behavior. Observed use of language or hypothesized dispositions to respond, habits, and so on, may provide evidence as to the nature of this mental reality, but surely cannot constitute the actual subject matter of linguistics, if this is to be a serious discipline.

Second, Chomsky wanted to stress that behavior/performance is (as we saw in our discussion of natural language semantics and its relation to notions like truth and reference), an interaction effect. Knowledge of language guides/provides the basis for actual use, but does not completely determine use. It is only one of the contributing factors. Grammars interface with various cognitive components (and their signature limits) when language production and comprehension take place. We would, ideally, like to be able to identify the distinctive contribution of linguistic knowledge, but there is no ready-made recipe in this regard.

Consider the fact that a sentence like *I almost had my wallet stolen* is, as we saw in Chapter 7, three-way ambiguous. This must be part of our characterization of a speaker's knowledge, but given that it is so hard to imagine how one would attempt to steal one's own wallet, it is fair to say that this sentence will never be used with this meaning. This is but one case where use obscures knowledge.

Consider also the fact that our language faculty is designed to generate nested dependencies (recall the "anti-missile missile" example in Chapter 4). As Miller and Chomsky⁵ concluded in a seminal 1963 paper on the relation between competence and performance, such a system will generate sentences involving center embedding like *The cat that the rat saw ran*, as easily as tail-embedding sentences like *I caught the cat that the rat saw*. But center-embedding structures quickly degrade, as in *The rat that the cat that the dog chased caught died*, which everyone would agree is worse than *The cat that caught the rat that stole the cheese ran*. Miller and Chomsky demonstrated that it would complicate the theory of the language faculty dramatically if we imposed restrictions on how many levels of center embedding were possible (as we just saw, one level of embedding seems OK, but two seems too much). Miller and Chomsky concluded, quite reasonably, that the unacceptability of multiple center-embedding structures was a performance effect – it was due to a host of non-linguistic specific factors (memory limitations, the difficulty of carrying out an operation like center embedding while another operation of the same kind is being carried out, etc.). The fact that some multiple center-embedding structures sound much better than the one presented above, like *The reporter everyone I met trusts reported the coup*,⁶ suggests that this is the right conclusion. It shows that the competence–performance distinction is necessary to maintain a certain level of simplicity in the theory. Please note that the competence–performance distinction should not be used as a shield:⁷ one should not say "If my theory makes

the wrong prediction, I'll just say it's the right prediction, but it's obscured by performance factors." The competence–performance distinction does not insulate linguistic theory from some sorts of evidence ("I don't have to pay attention to this sort of experiment, after all, I study competence"), nor does it isolate cognitive neuroscience from the discoveries/advances of linguistic theory ("Why should I read this linguistic theory stuff? After all, I study performance.").

The competence–performance distinction is there to help linguists and other cognitive scientists cut the language faculty at its joints, and appreciate the fact that putting cognitive faculties to use is a very complex, interactive phenomenon. The distinction reminds me of a fact I learned in my math class in high school. My math teacher asked us to imagine a very big sheet of paper, longer, wider, but with the same thickness as our familiar sheets of paper. He then asked us how many times we would have to fold this piece of paper to arrive at a thickness of 1 meter, or a thickness corresponding to the distance between the Earth and the moon. I still remember being shocked when I learned that it would take a ridiculously small amount of foldings to get there (I encourage the reader to try to come up with the answers). But on second thought, it's easy to see why very few steps would carry you very far: each time we fold, we double the thickness of the paper. We are dealing with a case of exponential growth, and our intuitions in this domain are very shaky; it's one of those mental tunnels Massimo Piattelli-Palmarini⁸ has eloquently written about. But the point of this example is that after we arrived at the right mathematical formula, my teacher told us that this is just a thought experiment. Due to the physical force of resistance, we will never be able to physically fold the sheet of paper the amount of times required to get our results. In hindsight, this is another illustration of a competence–performance distinction. On the one hand we have a very simple theory, a predictive formula (competence); on the other we have physical forces that prevent us from realizing some of the predictions of the formula (performance effects). Obviously, it would be wrong to complicate the formula to track the influence of (independent) physical forces. For the same reason it would be wrong to complicate our theory of competence to track the interactions underlying performance. (Remember it would complicate the theory considerably if we tried to capture what we know about our language, while taking into account the fact that no infinite sentence has ever been produced.)

The third thing that Chomsky wanted to highlight with the competence–performance distinction is this: it would be foolish to restrict the study of linguistics to the specific linguistic productions of speakers. I am emphasizing this because the perceived objectivity of corpus-based studies remains a fatal attraction to many students of language. The study of linguistic knowledge should not only take into consideration what speakers produce and perceive; it should also capture what speakers can't do, the "negative facts" we talked about in Chapter 7. Identifying how speakers can't use language (which, by definition, will never be reflected directly in a corpus) reveals the constraints that Universal Grammar imposes. In other words, what speakers systematically avoid doing (the sentences they always reject

as unacceptable, the sentences structures they never construct) is as rich a source of data as what speakers do (it may even be a privileged source of data). One should never lose sight of the fact that the human language capacity extends far beyond what one finds in corpora.

2 Putting Grammar to Use: Prospects and Problems

With these important considerations in place, we can now turn to attempts to integrate models of linguistic knowledge into models of comprehension, production, and judgments – the three main sources of data that linguists (and other cognitive scientists) rely on to infer knowledge from behavior. Examining how well hypotheses about knowledge of language fare when confronted with data of various sorts will enable me to show that a great deal of evidence from various domains converges to support theoretical hypotheses currently being entertained. The fact that data from acceptability judgments, sentence production, and sentence comprehension all point to the same conclusions is a very good thing, since it means that the empirical basis of linguistic theory is quite solid, drawing from a very large body of evidence (larger perhaps than some might think). I will also show that although competence and performance should be kept separate, there is a remarkable degree of transparency between (our theories of) knowledge of language and (our theories of) language production/comprehension. That is, the real-time deployment of grammatical knowledge is highly accurate. As some researchers⁹ have noted, this degree of transparency and accuracy could have a very desirable empirical payoff: if grammar is fairly transparent in real-time comprehension/production, it means that real-time brain activity measurements could be used to detect the anatomical and physiological “wetware” support of our mental structures. In other words, it would make it possible for all sorts of experimentation techniques to bear on the nature of linguistic representation as much as data from acquisition does. It also suggests that linguistic theory can be held responsible for the body of phenomena revealed even by neurological approaches to language.

The degree of convergence of evidence, and the transparency just mentioned, show that the hypotheses entertained by what one might call “armchair linguists” are to be taken very seriously; indeed, they suggest that such hypotheses may serve as the basis for a cognitive neuroscience of language, which many see as the “holy grail” of neuroscience, and about which I will have more to say in Part IV.

The point just made is, I think, a very important one. The original conception of linguistic theory (“generative grammar”) suggested that it could serve as a sound foundation for the integration of neuroscience with cognitive science. Although it is often thought that linguists, studying competence, deal with native speakers’ intuitions while cognitive neuroscientists build theories of language that make predictions about reaction times and brain activations, the divide is not real.

As Chomsky made clear in the first chapter of *Aspects of the Theory of Syntax*,¹⁰

The study of performance models incorporating generative grammars may be a fruitful study; furthermore, it is difficult to imagine any other basis on which a theory of performance might develop.

Cognitive science and cognitive neuroscience are continuous enterprises; we can and should seek to integrate them further, but one is not reducible to the other.¹¹ Linguistic theory does not study a Platonic object, an abstract object about which speakers may have (stable) intuitions. Behavioral data of all kinds, including language production and comprehension in real time, are important, and certainly don't belong to an independent field distinct from linguistic theory. As the quote from Chomsky just reproduced makes clear, there cannot be a totally independent field of psycho- or neuro-linguistics. As for any experimental science, as Alec Marantz¹² points out, predictions of experimental results and constructions of experimental designs require at least a rudimentary theory of the task in the experiment; they require some familiarity with the results of linguistic theory. Obviously, the stronger the familiarity, the better the experiments.

It is true, though, that after a decade of productive interaction between theory and experiment in the 1960s (epitomized by the joint works of Miller and Chomsky),¹³ where the new theoretical hypotheses were the strongest guiding force behind the experiments,¹⁴ theoretical and experimental investigations grew apart in the 1970s. Part of this is inevitable. Language use is, as we saw in the context of meaning in Chapter 7, what one might call an interaction effect. As the field of cognitive science matured, as people learned more about the various factors involved in language use, the contribution of linguistic knowledge came to be seen, quite appropriately, as one of several factors at work, and, as a result, the role of theoretical investigations has come to occupy a less central place than it did when it was virtually the only thing people knew about.

But part of the dissociation was the result of a perceived failure, reported by Fodor, Bever, and Garrett in 1974¹⁵ and always mentioned in any discussion of the performance–competence relation. The failure concerned the so-called “derivational theory of complexity” – the hypothesis, guiding Miller and Chomsky's joint work,¹⁶ that the number of operations that the grammar requires to generate a structure according to linguistic theory should correlate with, e.g., the reaction time of speakers processing the sentence in some psycholinguistic task. That is, in the words of Alec Marantz,¹⁷ all other things being equal (and, of course, performance being an interaction effect, this caveat is crucial), the more complex a representation – the longer and more complex the linguistic computations necessary to generate the representation – the longer it should take for a subject to perform any task involving the representation (and, although this was less of an issue in the 1960s, the more activity should be observed in the subject's brain in areas associated with creating and accessing the representation and with performing the task).

According to Fodor, Bever, and Garrett, the results of the experiments they were reporting on (about which I will have more to say below) suggested the following: although the linguistic representations argued for in linguistic theory were “psychologically real” (i.e., part of our knowledge of language), the mechanisms proposed by linguists at the time to create these representations/structures did not constitute the only way available to speakers for creating the representations of their language in real-time comprehension and production. Rather there might be separate strategies, to be studied in the field of psycholinguistics, for structure building that to some degree made it possible for speakers to bypass the application of rules proposed in linguistic theory. In Fodor, Bever, and Garrett’s words:¹⁸

The discovery that psycholinguistics has a subject matter – a body of phenomena which are systematic but not explicable within the constructs manipulated by formal linguistics – is, perhaps, the most important result of the last decade of psycholinguistic research.

A third reason, discussed by Marantz,¹⁹ why theory and experiment grew apart is because until very recently, the data of the sort commonly gathered by psycholinguists, based on reaction time in a variety of experimental situations, was not regarded as crucial in the day-to-day work of most linguists. As Marantz points out, linguists are more concerned with so-called static aspects of linguistic representations, the inventory of categories, and their structural relations. To address questions pertaining to these aspects, it is often sufficient to rely on native speakers’ intuitions (acceptability judgments). In other words, there was (and, to some extent, still is) enough data to keep linguists busy without having to resort to complex and costly experiments. So part of the reason why theory and experiment grew apart was pragmatic: “cheap” data is available, so why should one look for something else?

3 A Note on Acceptability Judgments

This last point deserves a little digression, though. When I say that theory and experiment grew apart, it is important to stress that although they are rarely reported/presented as such in publications, the native speaker’s intuitions on which theoretical linguists rely are (very cheap) experiments.²⁰ Although they are not often presented as the results of a behavioral experiment, judgments are behavioral data. All the examples mentioned in previous chapters to illustrate this or that result concerning our linguistic knowledge are the results of experiments where a linguist asked a native speaker (sometimes herself) how good a sentence sounds compared to another, what sort of interpretation this or that example supports, etc. Although the term “grammaticality judgment” is often used for this kind of mini-experiment in the literature, it is a misnomer. Speakers lack intuitions about grammaticality.

Asking about whether a given sentence is grammatical or not would be asking whether the informant's mental grammar establishes a particular sound–meaning pair. No one can have intuitions about that, any more than one can have intuitions about the nature of physical forces or the arrangement of molecules. Intuitions are about acceptability (hence the more appropriate label “acceptability judgment”). And although no one doubts that judging the acceptability of sentences is an interaction effect, with all sorts of factors involved, it turns out that speakers, when freed from time constraints, have very stable judgments about a large class of sentences, a stability that has served up to now as a solid basis for linguistic theory.²¹ By presenting speakers with lists of sentences often involving minor variations on a theme (“minimal pairs”; i.e., minimally different examples), linguists have tried to control for various factors that might influence speakers' judgments, and from the results of these informal experiments, they have inferred what we take to be our best guess regarding the nature of the language faculty.

But although judgments will remain part and parcel of a linguist's toolbox for probing Universal Grammar, theoretical linguists would readily admit that they are no panacea. As linguistic theory deepens, it demands that increasingly subtle facts be tested. There is no guarantee that cheap experimentation with acceptability judgments will be sufficient.

As a matter of fact, we already know of several cases where judgments are less robust than one might want. In such cases, one should perhaps let the theory decide which judgments reflect the true grammaticality of the structure, but in some other cases, where the theoretical predictions are not so clear, it may be necessary to rely on a bigger pool of informants, or use better samples (with better controls).²² As in any empirical science, the construction of the database is ongoing work, and it may well be necessary to explore different methods to triangulate the phenomena.

Recently, the role of judgments has come under attack. A series of studies²³ have emphasized the gradient character of acceptability, i.e., the existence of many shades between totally acceptable and totally unacceptable. Many have since argued that gradience should be built into our theory of knowledge of language, but this may be too quick a move. One should not lose sight of the fact that judgments are interaction effects. Gradient acceptability may be the consequence of the subject of the experiment trying to fit too many different properties on a single scale (from acceptable to unacceptable). As Colin Phillips²⁴ has pointed out, gradience may reflect the combined response to the grammar's ability to generate the sentence, to the violation of grammatical constraints, to the possibility (or ease) of recovering the intended meaning, to the salience of some reading, to the availability of alternative ways of expressing the same meaning (that the speaker can think of during the experiment), and, no doubt, a host of other factors. I agree with Phillips that it is clearly not enough to document gradience; one must search for its cause(s), and no doubt this will be a big part of the immediate future of linguistic inquiry. To identify the causes of gradience, it may be necessary to draw on multiple techniques to complement what we have already discovered on the basis of judgments, or supplement what we can't discover through judgments.

4 Grammatical Transparency: Six Arguments

It so happens that in the past few years, the realization that new techniques should be tried to probe the nature of our language faculty has gone hand in hand with the realization that the reported failure of the derivational theory of complexity may have been premature.²⁵ Since the derivational theory of complexity would serve us well in trying to interpret data from reaction times or brain imaging (online measurements) to complement judgments (offline measurements), it is, I think, helpful to re-examine the body of phenomena that Fodor, Bever, and Garrett (and many subsequent studies) deemed outside the range of constructs manipulated by theoretical linguists.

There are a number of arguments that are standardly used to argue that there is a discrepancy between what we know about language and how we parse or produce sentences, leading to the claim that there are strategies that may be specific to language, but not part of our knowledge of language, that determine how we understand and produce sentences. Many arguments go back to the 1970s, and have rarely been revisited since, except by Colin Phillips,²⁶ upon whose reflections I will build here, and whose conclusions I endorse. The list of arguments includes:²⁷

- 1 the constructs assumed in linguistic theory do not appear to be suitable for direct deployment in speaking or understanding
- 2 parsing and production lack the precision required of, and associated with, grammatical knowledge
- 3 grammars typically do not provide the tools needed to account for well-established parsing phenomena such as “garden-path” sentences (on which, see below)
- 4 the apparently slow and effortful nature of acceptability judgments suggests the existence of a system that operates on a different time scale from parsing and production, which are extremely rapid
- 5 speaking and understanding are clearly different things, which break down in different ways; they must therefore be the products of different systems
- 6 the derivational theory of complexity was a good, strong hypothesis, arguably the null hypothesis, or the most natural way to understand the relationship between the grammar and the parser/producer, but many studies have shown it to be inadequate

Each of these arguments has been constructed over series of experiments, and taken together they are often assumed to be decisive in thinking about how knowledge is put to use. However, as Phillips has shown, careful examination of each of these arguments reveals that they are not as compelling as one might think, and that a return to the derivational theory of complexity may be fruitful.²⁸

Let me summarize Phillips’ points here, going through each argument in the order I have listed them above.

(1) *Grammars don't look like parsers*

As I discussed at the end of Chapter 4, linguists currently think that the best way to understand how mental structures that are part of our linguistic knowledge are built appears to be an algorithm that takes elements and merges them two by two, recombining elements as the construction proceeds. Visually speaking, the algorithm proceeds from the bottom of a tree representation all the way to the topmost node, a bit like the way houses of cards are built. Accordingly, the structure of a sentence like *John will kiss Mary* emerges from first merging *kiss* and *Mary*, then merging the resulting set (“Verb Phrase”) with *will*, and the resulting set to *John*. The end result is a sentence made up of constituents (phrases) that reflect the history of the derivation: [John [will [kiss Mary]]].

This is depicted as follows:

- a. {kiss Mary}
- b. {will, {a}}
- c. {John, {b}}

The constituents that are formed by this bottom-up procedure appear to be exactly what are needed to capture a wide range of facts at the heart of linguistic theory. However, it is clear that we do not parse sentences that way; we do not concatenate words beginning with the last two words of a sentence; instead we parse (combine) words from left to right, in the order received from the source of information. We are thus dealing with a difference in the way information flows (bottom-up or left-to-right).

This suggests a natural (though perhaps incorrect) hypothesis:²⁹ the applicability of grammatical principles must be reversible (i.e., they must be directionally invariant). (Exactly the same situation obtains in theoretical physics, where equations like those of classical/Newtonian mechanics are time-reversible, even if “real” time flows only in one direction.)

Hearers clearly do not wait until the end of a sentence to begin interpreting groups of words. There is massive evidence³⁰ that they seek to narrow down the interpretive paths as they parse, which suggests that they build partial constituents (“treelets”) at each stage, which they fill in as they proceed, assigning to them as much interpretation as they can (reanalyzing/revising the interpretive guesses as they encounter more words/build more constituents). In other words, speakers seem to be able to build sentence skeletons on the basis of meager evidence (a few words is enough to guide them, especially in what structures should not be entertained); that is, their knowledge of grammar seems to help them use already-seen material to project not-yet-encountered material.

Needless to say, the space of possibilities that the hearer entertains as she parses the sentence is not only constrained by the projected skeleton made possible by her grammatical knowledge; other factors having to do with guessing the speaker’s intention and so on also play a role. In fact, it is clear that speakers unconsciously use all the possible tools at their disposal to succeed in parsing utterances.

I agree with Phillips that the difference in flow of information need not be seen as an argument against the real-time deployment/accessibility of linguistic knowledge, as speakers appear to be able to build structure incrementally and accurately, even when presented with sequences of words that cannot be combined directly using the rules of grammar.

(2) *Parsers don't look like grammars*

This argument is the mirror image of the one just discussed. If you have ever looked at an introductory discussion of sentence comprehension,³¹ you are bound to encounter studies of comprehension breakdown, such as the multiple center-embedding examples mentioned above, or the famous garden-path sentences, such as:³²

The horse raced past the barn fell.

where speakers gag once they hit upon *fell*, as they tend to interpret *The horse raced past the barn* as a complete clause, not as a noun phrase containing a dependent clause: "The horse (that was) raced past the barn . . ." (as in the minimally different *The horse ridden past the barn fell*, which poses no comprehension breakdown).

The argument for a parsing system quite different from the one made available by the grammar is that grammar is, by definition, accurate and precise (albeit quite slow), whereas parsing is fast, but prone to errors. However, I agree with Phillips that in order for the argument to be convincing, it would need to be shown that hearers construct hypotheses that go against what their knowledge of grammar allows. The language faculty clearly allows structures corresponding to "the horse raced past the barn." As Phillips writes,³³

garden path sentences arise in circumstances of structural ambiguity, where two or more possible structural (grammatically licit) analyses are available. If the wrong choice is made during parsing, and subsequently breaks down when it becomes clear that the choice was the wrong one, this reflects lack of telepathy, not lack of grammatical precision.

In no case does the listener appear to construct hypotheses that are outside the realm of what her grammar allows. As for center-embedding examples, the consensus view since Miller and Chomsky's original study³⁴ is that these are the reflex of memory overload, not grammatical imprecision.

Perhaps the most serious challenge to the claim that hearers don't entertain hypotheses that the grammar rules out comes from sentences discussed by Phillips like:³⁵

While the man hunted the deer ran into the woods.

Most speakers go down the garden path here, as they interpret *the deer* as the object of *hunted* before realizing that this was the wrong choice. What is most interesting about this example is that even after recovering from this effect, it is reported that the subjects of experiments continue to believe that the deer was hunted. It is as if they treat *the deer* as both the object of *hunted* and the subject of *ran*, which is grammatically impossible (the only possible structure is: [while the man hunted] [the deer ran into the woods]; the “shared” structure [while the man hunted [the deer] ran into the woods] is not something Universal Grammar allows). Phillips³⁶ suggests that such cases indicate that some previously entertained hypotheses tend to linger longer than they should, and that it takes more than structural reanalysis to completely dismantle options once deemed plausible.

Be that as it may, it is important to distinguish grammatical breakdown/failure from breakdown/failure in general. What is required of the grammar is that it captures possibilities entertained at any given time, not that it predicts which option will turn out to be the correct one.

(3) *Grammars fail to explain ambiguity resolution phenomena*

In the wake of the demise of the derivational theory of complexity in the 1970s psycholinguists proposed a series of parsing-specific strategies, heuristics, and biases that they felt were needed in light of the perceived failure of grammatical support. I personally don't doubt that speakers form habits and biases that account for some parsing preferences, or lead to the storing of (parts of) utterances that fossilize as idiomatic expressions.³⁷ But these in no way replace grammatical knowledge in guiding the production and interpretation of sentences. As a matter of fact, some of the strategies once proposed as parsing-specific mechanisms turned out to correspond to principles of grammatical knowledge that were subsequently discovered. For example, it is well known that speakers hearing a question word at the beginning of a sentence try to interpret it in the first possible place from which it could have moved (this is often described as a “filler” looking for a “gap”). So, when subjects are presented with a reading task where they are shown one word at a time and have to press a button to advance to the next word, they slow down once they realize that they interpreted the question too quickly, as in *What did you say . . . that Bill read?* Speakers slow down when they hit upon *that* because their first guess was that the complete question would be *What did you say?* But in this case interpreting the question word at the first available opportunity (a reasonable parsing heuristic of the sort that was proposed in the 1970s) has the same effect as the grammatical principle that says that transitive verbs like *say* require an object.³⁸ The fact that *what* is interpreted in conjunction with *say* may not have to do with an attempt to interpret the question word as soon as possible, but may in fact be a reflex of an attempt to provide the verb with the very first element available to be its object, in accordance with the principles of Universal Grammar.

The point of this discussion is to show that what one may have thought was a genuine parsing-specific strategy may turn out to have the same effect as a principle

of grammar subsequently discovered. One should bear in mind that the body of doctrine made available by theoretical linguists is still very much work in progress. The fact that 1970s-style grammatical theory failed to account for the way sentences are parsed does not mean that no grammatical theory will ever be able to do so. (In fact, my impression is that progress in theoretical linguistics renders the real-time deployment of grammatical knowledge more and more plausible, but I won't be able to substantiate this point here, as it would require an in-depth discussion of current linguistic theory.)

(4) *Parsing and production are fast, acceptability judgments are slow*

Many acceptability judgments, especially those involving subtle semantic contrasts, are slow and difficult. This has been taken as a sign that the grammar operates on a different time scale from real-time parsing and production. But the evidence does not support such a strong conclusion. Several experiments point to the fact that speakers can detect grammatical violations within a few hundred milliseconds of the presentation of the offending word. As for instances of slow (but accurate) judgments, they may reflect repeated attempts to re-parse the sentence, perhaps to avoid an irrelevant initial parse, or construct an appropriate mental scenario to make additional reading plausible (consider again the example *I almost had my wallet stolen*, where speakers are asked to detect various readings).

More troubling are cases where sentences that should be judged unacceptable (because they are ungrammatical) are first judged acceptable, as in the infamous *More people have been to Russia than I have*.³⁹ Upon reflection, speakers come to realize that such sentences are nonsensical (this is clear in *More people have been to Russia than I have been*), but at first they sound just as good as minimally different ones, like *John has read more books than I have*, or *More people have been to Russia than I could imagine*, or *Many people have been to Russia more often than I have been*. Although a lot about these misleading sentences remain to be understood, I think it helps to remember that judging a sentence is also an interaction effect. It may well be that the grammar rules out the sentence as it is parsed, but other factors (overall similarity among sentences already encountered, etc.) may well boost the acceptability of some sentences.

(5) *Speaking and understanding are different*

Traditionally, production has been thought to track grammatical knowledge much more faithfully than parsing; not surprisingly, then, classic models of parsing and production look quite different. Furthermore, the characteristic errors in parsing (garden paths) look quite different in nature from the characteristic errors in production (slips of the tongue, spoonerisms).⁴⁰ Also, research on language disorders has revealed numerous cases where comprehension and production are affected differently (see next chapter).

However, differences between the outputs of parsing and production need not lead to the conclusion that independent systems are involved. A single sentence structure generator that pairs sound/sign and meaning may underlie both processes, just

as the hand may be put to use in many different ways, but it is the same biological endowment underlying all its uses. As in the case of argument (1), differences in the flow of information need not lead to the claim that knowledge fails to guide the various uses to which it is put. As Phillips observes,⁴¹ it is interesting to note that in the area of word processing, the many differences between the tasks of picture naming and word recognition do not lead to the claim that we need two models of the lexicon, or different sets of word-formation strategies.

(6) *The alleged failure of the derivational theory of complexity*

As Marantz put it,⁴² the derivational theory of complexity is the null hypothesis for the working linguist. All else being equal, the character of the grammar should be reflected in its use. Short of that, linguistics may take on a Platonist character, which it lacks once conceived of as a field of cognitive science (as I have done so far, following Chomsky). I agree with Marantz that the derivational theory of complexity demystifies the nature of linguistic representations and computations.

The derivational theory of complexity failed in the 1970s, but this may simply have been due to the fact that we didn't have the right details for linguistic theory. Miller and Chomsky were right to suggest a close (transparent) relationship between the mechanisms of grammar and real-time sentence processing, but they may have been wrong about the specific mechanisms of the grammar. Theirs was a very specific attempt to link knowledge and use, and its failure need not have led to the wholesale rejection of the attempt to recognize the application of grammatical principles in real-time processing.

5 Summary

All in all, a growing number of experiments carried out by researchers who know about experimental techniques and about the details of linguistic theory point to the fact that real-time language processes show great grammatical precision and faithfulness. As I already pointed out, adopting the derivational theory of complexity as a working hypothesis widens the domain of data relevant for grammatical studies; the more data we have, the more falsifiable the theory, the better the science.

If, instead of assuming, like so many psycholinguists still do, that there exist specific mechanisms to provide a rough and ready approximation of what the grammar allows, and that the details of grammatical knowledge come in much later than the first parse, we instead put the findings of theoretical linguistics at the core of what we do with language, the field of linguistics and of cognitive science as a whole will have a much more integrated, unified character.

Let me stress that the mentalist commitments discussed in this chapter, though they have been discussed in the context of the study of the language faculty, are valid for all cognitive faculties. Be it music, morality, or arithmetic, any body of

knowledge put into use requires one to be clear about the competence–performance distinction, and how behavioral data may shed light on the cognitive capacity underlying behavior. It stands to reason that “cheap” experiments, whenever possible, will form the starting point of investigation, but experimental evidence of any kind should be welcome, as we seek convergence and unification.⁴³

PART IV

Missing Links

The (Mis)Measure of Mind

1 Connecting Mind and Brain

In the last part of this book I would like to address questions that are at the limit of our understanding – questions that get to the heart of what some have qualified as the “hardest problem in science,”¹ or the “holy grail of cognitive neuroscience,”² touching on “the most complex object in the universe”³; and, not surprisingly, questions “we will never answer.”⁴ The questions revolve around the relation between the mind and the brain, and the evolution of mental faculties. As I have done throughout this book, I will use language as a case study, but the issues addressed in the following pages arise for many other mental organs, although because we have a fairly detailed picture of the mental structures involved in language, some of the difficulties to be discussed come into sharper focus in this area. This is both a good thing and a bad thing. A bad thing because the gulf between neuroscience and cognitive science is vast, and people prone to pessimism may lose hope. A good thing because once the nature of the difficulties is clear, we can take a few steps toward redirecting inquiry toward possible solutions. Although I sympathize with the skeptics, I am enough of an optimist to believe that it’s worth trying to see how far we can go when taking the results and methods of theoretical linguistics very seriously.

When discussing issues such as how mental faculties may be implemented in the brain, or how they many have evolved, it is too easy to begin and end each paragraph with “we don’t know (yet).” We know that somehow the mind emerges from the brain (the mind is what the brain does, in one popular statement),⁵ and we also know that somehow the brain has evolved in the course of history to produce these mental organs. But this is not news. Although it is sometimes presented⁶ as an “astonishing hypothesis,”⁷ a “radical new idea,”⁸ or a “bold assertion,”⁹ the claim that mental phenomena are entirely natural and caused by the neurophysiological activities of the brain was familiar to the natural philosophers of the seventeenth and

eighteenth centuries,¹⁰ who had already concluded that properties “termed mental” reduce somehow to “the organical structure of the brain” (Joseph Priestley),¹¹ and were comfortable with the idea that thought is a “secretion” of the brain.¹² As for the idea that mental faculties have evolved, it’s been with us in a very explicit form since at least Darwin’s *Descent of Man*.¹³ Since then our understanding of biology and cognition has reached a very high degree of specificity; we have reasonably good descriptions of both genotypes and phenotypes but we are in dire need of “linking hypotheses.”¹⁴ When it comes to understanding how to relate genes and brain development, or brain circuitry and cognitive faculties, we are in the dark. We not only lack clear results, we also lack clear questions that would eventually allow us to get these results.

At the end of this part of the book the reader should not expect to find a way to make a seamless connection between mind, brain, and behavior, but I hope to be able to indicate a few ways in which this goal may look a bit more attainable.¹⁵ I will begin by outlining the classic model of the relation between language and the brain that finds its way into every textbook, and show (without denying its usefulness) how limited it is. Then I will turn to an outline of research questions that may prove more productive than the ones that are typically asked in the connection of mind and brain. In the next chapter, I show¹⁶ how the very same kind of questions may help shed light on issues pertaining to the evolution of mental faculties. In the last chapter of this part of the book, I reopen the issue of specificity of mental systems in light of the preceding discussion.

2 The Classic Model

Serious efforts to understand the neural substrate of language go back at least 150 years, although relevant observations go back much further (as far back as the Ancient Egyptians and the Greeks).¹⁷ Until the recent emergence of non-invasive brain activity recording techniques such as positron emission tomography (PET) and functional magnetic resonance imaging (fMRI), interpreting findings from brain-damage patients was virtually the only way to infer how language may be implemented in “wetware.” Such investigations have led to the Broca-Wernicke-Lichtheim-Geschwind, or classic, model, the first large-scale functional anatomical proposal on language processing, and one of the first comprehensive accounts of a higher cerebral function. It was the result of cumulative efforts by Paul Broca,¹⁸ Carl Wernicke,¹⁹ and Ludwig Lichtheim²⁰ during the nineteenth century, and it was revived and modified some 40 years ago by Norman Geschwind,²¹ after it had been forgotten for about half a century.²²

Broca’s main discovery, originally reported in 1861,²³ is perhaps the best known of all. Broca discovered a French patient (Mr Leborgne) who had been unable to pronounce anything but a single syllable (“tan”) (hence the patient’s nickname “Tan”) after reporting severe pain in the brain. A postmortem analysis revealed a large lesion

in his left inferior frontal cortex (an area just forward from and slightly above the left ear). Based on this correlation, Broca concluded that this area had a necessary role in speech production (following a similar suggestion made earlier by Jean-Baptiste Bouillaud). Thus was born Broca's area. A few years later, Wernicke²⁴ observed that lesions in another area further back in the brain were correlated with fluent but nonsensical language production as well as impaired comprehension. Wernicke speculated that this part of the brain played a necessary role in language perception. Thus was born Wernicke's area. Lichtheim²⁵ synthesized Broca's and Wernicke's claims and posited a connecting area (involving several regions) between Broca's area and Wernicke's area, assumed to be responsible for semantic processing. Geschwind²⁶ adopted most of Lichtheim's assumptions, but modified his claims about semantic processing (in a way that I won't review here, as it does not bear on the main focus of this chapter).




Figure 10.1 Broca's area and Wernicke's area

The emerging picture is the one represented in Figure 10.1, with Broca's area (inferior frontal area), Wernicke's area (superior temporal region), and a connecting fiber tract (arcuate fasciculus). In the context of Brodmann's map of cortical areas (areas identified and numbered in the 1908 report by Korbinian Brodmann²⁷ based on the organization of neurons he observed in the cortex), Broca's area corresponds to area 44/45, and Wernicke's area to area 22, as depicted in Figure 10.2 on p. 152.

The classic model, as its name suggests, has remained at the core of most textbook discussions of the neural basis of language. As David Poeppel and Greg Hickok point out,²⁸ one major feature of the classic model is that it was the first of its kind. Its historical significance should not be underestimated: it was one of the first major attempts in cognitive neuroscience (or neuropsychology), laying the foundation for the principle of functional localization that still dominates the field. Another,




Figure 10.2 Brodmann's map of cortical areas (from Phillips and Sakai 2005)

more substantive advantage listed by Poeppel and Hickok²⁹ is that its main predictions still serve as very useful guides for classifying both lesions and syndromes, despite the numerous inevitable extensions and modifications to the theory. To this day, clinical practice is guided by this classic model, and it has played a key role in the treatment of aphasia (language disorders).

When one goes back to, for example, Wernicke's writings, as I did to write this chapter, and as Poeppel and Hickok³⁰ did to write their overview, one cannot help but be struck by how prescient some of his reflections were. The goal of the model was very clearly stated: it was to develop a psychological model of language on an anatomical basis, which remains our goal today, and one recognizes in it the seeds of many of the leading ideas of contemporary research on the neural basis of speech perception and language comprehension. To be sure, we now have different tools at our disposal, and we have refined some of the concepts that were at the heart of the classic model, but many of the central issues there remain at the forefront of research now. Perhaps the most significant refinement is this: Whereas the classic model of aphasia emphasized a division of language areas based on tasks (speaking, understanding), modern aphasia research suggests that it may be more appropriate to differentiate language areas on the basis of the type of information that they preferentially deal with, such as syntax, phonology, or semantics.³¹ But by and large it is fair to say that the space of possibilities has remained remarkably constant.

3 Major Flaws of the Classic Model

And yet, we have rather strong reasons, it seems to me, to believe that the standard model is deeply flawed. Although we still don't know very much (certainly not as much as we would like), I think we can say with some confidence that the classic

model fails at three levels:³² first, it fails to account for a number of facts regarding aphasia (properly seen as a complex of symptoms); second, the linguistic foundations of the model are hopelessly naive; third, the anatomical assertions have not held up in light of subsequent observations.

Regarding the first point, we now appreciate the fact that clinical aphasic syndromes are comprised of variable clusters of symptoms. For example, although the most obvious clinical symptom of Broca's aphasia is labored language production, careful studies in the 1970s and continuing to this day have revealed that Broca's patients also have comprehension difficulties,³³ particularly in situations where successful comprehension requires close attention to function words and inflectional morphemes. Thus, individuals with Broca's area often misunderstand who did what to whom in a passive sentence like *The man was kissed by the woman* (often understanding it as "The man kissed the woman"), so perhaps Broca's area has a task-independent role in syntactic processing that makes it important for speaking and understanding alike. Exactly what that role is remains a hotly debated issue, but facts of the sort just mentioned suggest a much more complex architecture than the classic model (down to its modern instantiations) predicts.

As for the second major flaw of the classic model, its naive assumptions about language organization, this should already be very clear to readers of this book. The original comprehension–production distinction certainly strikes us now as hopeless, but even the subsequent cuts in terms of syntax, phonology, and semantics are obviously much too coarse. I agree with David Poeppel's assessment, to be discussed in more detail below, that what has so far not been seriously considered is that such linguistic domains as syntax and semantics are themselves not monolithic, but have rich internal structure with numerous subcomponents and computational properties. It is indeed very problematic to assign a label like "syntax" to a brain area without taking into account the fact that linguists have demonstrated beyond reasonable doubt that syntax is subdivided into many parts, each with specific computational needs (cf. Chapter 4). It is clear that the coarse categorizations of language functions one finds in the literature are not specific enough to support the creation of linking hypotheses that have a fighting chance of being correct. (Because this point is very much connected to the issues discussed in previous parts of the book, I will come back to it with more illustrations below.)

Regarding the third flaw of the model, its anatomical inaccuracy, we now know³⁴ that (the typical symptoms associated with) Broca's aphasia need not be caused by damage to Broca's area, and the same is true about Wernicke's aphasia and its relation to Wernicke's area. Likewise, the classic model recognizes an overwhelming left-hemisphere dominance for language, but arguments exist in the literature to the effect that some aspects of linguistic function, such as speech perception, are organized bilaterally rather than unilaterally in the left hemisphere,³⁵ and there is evidence that the so-called speech regions are not anatomically or functionally homogeneous. There is also evidence for plasticity. Although left-hemisphere dominance is the normal default pattern, we know that children who have undergone removal of the left hemisphere early in life to control intractable epilepsy show fairly good recovery of language abilities, which indicates that the right hemisphere

is able to take over many language functions if the left hemisphere is removed. So, all in all, there is a growing consensus³⁶ that we need to pay more attention to cortical areas outside of the traditional language zones, and that the right hemisphere must not be underestimated.

4 Why Then Is the Classic Model Still Mainstream?

At this point, the reader may well be entitled to ask why, if many of the neuropsychological, anatomical, and linguistic assumptions implicit or explicit in the classic model are known to be problematic, has the model enjoyed such a prominent role in research for so long?

The answer could be simply that it's the best thing we have, disappointing as it may be. But following important reflections by David Poeppel³⁷ on this issue, I think that this is only part of the answer. I believe that the reason the classic model has stood for so long is because of the sort of questions and guiding hypotheses people have been willing to consider. To gain a bit more perspective, it may be useful to consider the root of the classic model, and its cartographic tendencies (mapping functions onto brain regions).

As I already pointed out in the context of modularity in Chapter 8, the idea that different portions of the brain are associated with different processes and representations has a long history, going at least as far back as Franz Josef Gall (1758–1828), and his infamous phrenology. Gall's phrenology was crucially based on the idea that the human mind is made up of a list of mental faculties, each of which is tied to a certain location, as can be seen from Figure 10.3 on p. 155, a variant of which many readers may have encountered already.

We now view this approach as hopelessly naive and psychologically wrong – yet something like it is at the basis of most of the experiments involving brain imaging and, (very often, implicitly) provides the basic components of explanations in cognitive neuroscience. Yes, one can say that to this day, the questions asked in neuroscience have been unashamedly phrenological, so much so that one can speak, as Uttal³⁸ does, of a new phrenology.

As David Poeppel³⁹ has pointed out, there is a good reason why localization of function, and particularly the idea of what is called a topographic (spatial) map of the brain, has had such a lasting influence. In sensory and motor domains, the idea of a localized spatial map is actually correct (as far as we can tell). Research on the organization of the visual system has provided strong evidence that the brain's representation of the world constructed from optical input rests on “retinotopic maps”: adjacent objects in the visual world are adjacently represented in visual brain areas.

Much like in the visual system, the auditory system makes use of a form of spatial mapping, in this case, mapping by frequency (what is called “tonotopic organization”). Thus, low-frequency sounds are represented at one end of the




Figure 10.3 Gall's phrenology

map and high-frequency sounds at the other end. Any given sound is decomposed into its frequency components, which in turn are represented in their respective locations in these maps. So one can say that it is a pattern of locations and their activation that form the brain's representation of that sound.

A third domain where talk of maps in neuroscience makes a lot of sense comes from the representation of body surface in the brain, known as "somatotopy," or, as it is often represented in cartoon-like fashion, the somatosensory homunculus (Figure 10.4 on p. 156). As depicted, the body surface is represented in an orderly way in the parietal cortex. As the cartoon representation makes clear, some areas of the body are dramatically over-represented (magnified) in the somatosensory cortex. (The brain space allocated to the mouth, face, and hands is particularly large compared to other areas.) The mind/brain's view of the body is thus a little human- (homunculus-) shaped form with wildly exaggerated hands and head. That is to say, the body we feel, like the world we see and the sounds we hear, is represented as a spatial map.

The evidence behind the principles of retinotopy, tonotopy, and somatotopy is very strong, stronger than virtually everything else we know about the brain and its organization. It's been replicated using many techniques, and found valid for many species. As Poeppel⁴⁰ points out, given this success, it is only natural to try to extend




Figure 10.4 The somatosensory homunculus (after Penfield and Rasmussen 1950)

the notion of spatial topographic representation to domains *beyond* sensory and motor function. This may account, in part, for the cartographic tendencies in the field. Furthermore, cartography has led to some very interesting results, which I do not wish to undermine in any way. Such results have been made possible by the advent of non-invasive brain activity imaging technologies,⁴¹ which are of two sorts: those measuring blood flow (hemodynamic: Positron Emission Tomography (PET), functional Magnetic Resonance Imaging (fMRI)), and those relying on electric activity (electroencephalography (EEG), magnetoencephalography (MEG)).

Thanks to such techniques, it is now possible to determine a patient's lesion site shortly after damage occurs, rather than waiting for autopsy. This, as one may expect, has led to a dramatic increase in the database of knowledge available for deficit-lesion correlations. Also, modern techniques such as fMRI can be used to test for correlations between selective activation patterns in normal adults and selective deficits in patients. Some other techniques, such as TMS (transcranial magnetic stimulation), have even made it possible to non-invasively apply stimulation to create momentary activation or impairment in highly specific cortical regions to test whether a particular area is essential for a specific task, rather than merely involved in that task.

In addition to findings about localization of language in the brain, techniques like EEG and MEG measure the scalp voltages or magnetic fields generated by electrical activity in the brain, and provide a detailed record of the temporal dynamics of brain activity. As many⁴² have argued, studies relying on such methods may provide insights into the mechanisms that allow cognitive tasks like language processing to be so fast and efficient.

Brain imaging techniques have also reinforced one of the most strongly held beliefs among linguists: the idea that our biology delineates the set of possible languages (thereby excluding some systems that one would regard as logically plausible), the very idea of Universal Grammar.⁴³

In a series of experiments,⁴⁴ Andrea Moro and his colleagues asked Italian subjects to learn a small vocabulary in a foreign language that they had never learned at any point (Japanese in this case). After that subjects were taught rules of that language, but with a twist: the experimenters taught them two sets of rules (without telling them), one set that is compatible with what linguists hypothesize about Universal Grammar, specifically rules that are the result of parametric values that Italian does not make use of but that Japanese does use (for example, put the main verb at the right edge of the Verb Phrase; cf. Chapter 5); and another set of rules that (to the best of our knowledge) no human language makes use of – what linguists would regard as UG-incompatible rules (e.g., always put the negation marker after the fourth word in the sentence). Once subjects had been taught these rules, they were asked to judge whether novel test sentences with which they were presented followed or violated these rules. Brain imaging was used to see their brains in action when they were performing such judgments. Remarkably, when judging items involving UG-consistent rules, Broca's area (to be more specific, a subdivision of Broca's region; Brodmann's area 45, known as *parstriangularis*) was activated throughout the task, whereas with UG-inconsistent rules, Broca's area was activated at first, but it quickly disengaged, and some different part of the brain got involved.

Some linguists, and other cognitive scientists, may remain unimpressed;⁴⁵ they may say, we knew it all along! After all, there is plenty of compelling evidence of the more traditional sort (cross-linguistic studies, language acquisition studies, etc.) that point to the existence of Universal Grammar. And I certainly don't want to give the reader the impression that brain recordings provide stronger evidence for the existence of Universal Grammar (nor do Moro and his colleagues). But as I already expressed in Chapter 9, the more evidence from different sources, the stronger the claim, and the better the science. So, data from certain areas of the brain can be very revealing.

5 It's Really Time to Look for Something Else

But in all fairness, even the wonderful experiments of the sort just discussed barely scratch the surface, and obscure many glaring problems. We still have no idea regarding the underlying question of what special properties of the human brain allow it to support language, and what the distinctive properties (at the neuronal level) of the language areas (if such exist) are.

Consider Broca's area, a darling of neuroscience, and the major player in Moro et al.'s experiment just discussed. In a different set of experiments,⁴⁶ Moro and his

colleagues presented Italian subjects with a made-up Italian version of Jabberwocky. Subjects were given four types of Jabberwocky sentences, and were asked to judge them (as acceptable or unacceptable): the normal ones (as baseline data), sentences violating Italian word order rules (e.g., a determiner following the “Jabberwocky” noun), sentences violating rules of Italian phonology (unpronounceable Italian-Jabberwocky), and finally sentences violating various rules of morphology/syntax such as violations of subject–verb agreement (again using Jabberwocky words). Interestingly, brain recordings of the subjects (carefully controlled for, in ways I won’t discuss here) show activation of Broca’s area whenever syntax (as opposed to, say phonology) is involved. On the basis of this, one may be tempted to conclude that Broca’s area is the seat of natural language syntax.

But now consider the fact that Broca’s area has been reported to be active in a number of linguistic tasks that are not syntactic,⁴⁷ for instance auditory lexical decision tasks, studies of minimal pairs in tone languages, and phonological/phonetic tasks such as the discrimination of rapid phonetic transitions or the processing of phoneme sequences as opposed to hummed notes. So the claim that Broca’s area is exclusively devoted to syntax is incorrect. Even worse, as I already hinted at when we were discussing the classic model, Broca’s area is active in a number of entirely non-linguistic tasks:⁴⁸ motor activation, motor imagery, and rhythmic perception; so Broca’s area does not even appear to be specialized for language.

Results of this sort highlight the difficulties faced by the new, cognitively informed phrenology. They should, in fact, come as no surprise once we take the following two facts into account. First fact: we know that Broca’s area, for example, is an area of high anatomical and functional complexity, with subregions that are distinct not only in their microanatomy, but also in their patterns of long-range connectivity.⁴⁹ Second fact: we know that syntax is far from a simplex, monolithic, unstructured computation, so claiming that Broca’s area is specialized for syntax combines two widely implausible claims: one about cognition, and one about neuroscience. (I should add that Moro and his colleagues, whose experiments I have used as illustrations because they strike me as nice instances of careful design under the guidance of a leader in theoretical linguistics, are extremely cautious when it comes to interpreting their data. I’m sure they have tapped into something real, but I think they would agree with me that we still don’t know exactly what that is.)

6 Can We Do Better?

David Poeppel⁵⁰ has articulated what I take to be the real problem leading to this disconcerting and, in many ways, depressing state of affairs. Aside from the recognition that cartography is just a first step, and that patterns of connectivity (among other things) must be taken into account, the problem is that what Poeppel calls the “alphabets” (also called the *primitives* or the *elementary constituents* or the *ontologies*, or the *parts lists*) of linguistics and neuroscience are profoundly distinct. He is

correct that, however preliminary we take our knowledge to be in both domains, there are no obvious links between the inventory of the fundamental elements and operations postulated in linguistics on the one hand and in the brain sciences on the other. At this point, as we saw in Chapter 4, the fundamental elements for language research include concepts like head and phrase, and fundamental operations on such concepts, like Merge, concatenation, and linearization. Basic neurobiological elements include structures such as dendrites, neurons, cortical columns, and operations such as long-term potentiation, oscillations, and synchronization.⁵¹

Poeppel notes that at this stage in our understanding it is totally unclear how one could formulate a connection between any of the members of these two sets. They seem completely disjoint. He goes on to identify two major causes for the total absence of linking hypotheses in this domain: what he calls the *granularity mismatch problem* (or, “neurolinguistics in practice”) and the *ontological incommensurability problem* (or, “neurolinguistics in principle”).

The granularity mismatch problem refers to the fact that linguistic and neurobiological studies of language operate with objects of different resolutions (granularities). In linguistics, extremely fine-grained distinctions and generalizations are made, referring to very detailed aspects of linguistic representation and computation (a quick look at any journal in theoretical linguistics will confirm this assessment). In contrast, neurobiological studies still work with much broader conceptual distinctions. For example, as we saw, numerous cognitive neuroscience studies seek to investigate the neural basis of “syntax” or “phonology.” But as I already pointed out, we know that these domains are not monolithic, and have many sublevels of representation (“structure”). So, until neurolinguists try to look for units that match what theoretical linguists hypothesize,⁵² the conundrum we are in will not go away.

The ontological incommensurability problem discussed by Poeppel refers to the possibility that the units of linguistic computation and units of neurobiological computation may be incommensurable. That is, even if neurolinguists achieve a higher degree of resolution (taking steps toward solving the granularity mismatch problem), it may well be that no linking hypothesis will be established. As Chomsky has often pointed out, basing his reflections on the relation between physics and chemistry,⁵³ for unification to be achieved, one may need a serious rethinking of the fundamental concepts in both fields. Exactly what this would involve is a topic for a chapter in its own right, and it is addressed in Chapter 12. All I will say for now is that it is likely to involve a rethinking of the issue of cognitive specificity.

To give the reader a more concrete idea of the research path to follow, let me briefly mention the most ambitious attempt at unification I know of: Ben Shalom and Poeppel⁵⁴ have proposed we describe the functional neuroanatomy of language processing in terms of the intersection of three different aspects of language processing (sound, form, and meaning, or phonology, syntax, and semantics) and three different types of computational operations underlying various aspects of language processing (found in all three domains): “memorizing” (learning of new forms and retrieval of stored forms), “analyzing” (decomposing into component

parts, accessing subparts of stored items, feature decomposition), and “synthesizing” (combining) processes.

According to Poeppel and Ben-Shalom, the temporal lobe deals principally with memorizing/storing lexical items, the parietal lobe with analyzing these items, and the frontal lobe with synthesizing their representations. Moreover, spatially, there is a superior to inferior gradient in each lobe, with phonological information being mapped in more superior/dorsal fields, syntactic information in the middle, and semantic representations more inferiorally/ventrally. As Poeppel and Ben Shalom would readily admit, the model is still very crude, but it is a step forward, for it places the emphasis on computational properties. We can only hope that more of such models will populate the pages of neuroscience journals.

I should perhaps reiterate at this point that the problems discussed above are, of course, not unique to research on brain and language. When we consider questions in visual perception, memory, attention, reasoning and decision making, or even higher-order experiential issues (such as free will or consciousness, which for some reason fascinate the public more than anything else), similar problems (or worse ones) emerge. As soon as the psychological domains in question give rise to articulated theories dealing with subtle aspects of the sort dealt with in linguistics, we are dealing with high-resolution concepts, which is in sharp contrast with the manner in which neurobiological research talks about these domains.

7 Hard Problems

To give the reader a feel for the challenges that await neuroscience (once it has shed its old phrenologist mantle), let me mention some of the most salient challenges that some have already identified:⁵⁵

- “The problem of 2”: Take a sentence like *The small book is on top of the big book*, and focus on the word *book* that appears twice in the same linguistic event (utterance). Presumably, the word *book* is stored somehow in the brain. We may not know quite how, but even once we get the answer, we will still have to worry about what it means for the brain to maintain two instances of the same lexical item active and distinct in the same event. As Phillips and Lau⁵⁶ note, this requires the ability to perform the natural equivalent of copying the contents of an item in long-term memory into multiple registers in a computational workspace (this is actually what linguists do when they say that the word *book* is copied twice from the mental lexicon in the course of linguistic computation),⁵⁷ but what is the equivalent of a computer register in neural terms?
- The problem of rules containing variables: As we discussed already, our knowledge of language makes it clear that what we know must be stored intensionally, not extensionally: no finite brain could contain the infinite range of expressions our language faculty enables us to construct. What must somehow be stored in the

brain is a set of rules containing variables: statements like “For each Verb x , combine it with at least one Noun y , and call the resulting set w ,” or “All phrases must have a head” (cf. Chapter 4). How the brain does this is completely unclear. The problem is particularly acute⁵⁸ when we realize that in addition to these abstract structural templates and rules, the brain must also be very good at capturing the fact that some very specific linguistic combinations must be memorized as well (as in the case of idiomatic expressions).

- The binding problem: Consider the word *book* again. It’s a very rich piece of language, about which linguists know a lot: it has a specific phonology (sound structure), it belongs to a specific category (noun, which we must keep separate from the verb *to book*), and a rich semantics (recall the abstract/concrete interpretations of *The book that Joyce wrote was too heavy for me to carry out of the library* discussed in Chapter 7). All that information, which may well be distributed all over the brain (as many brain imaging techniques suggest), must come together when we process the word *book*. How this binding together of information takes place is unknown, and this says nothing about the binding together of information of multiple words, which must take place every time we process a sentence.
- The discreteness problem: Linguistic models deal with discrete categories (computational units), but most well-known phenomena at the neurophysiological level appear to be continuously varying, non-discrete (analog) phenomena (as Marantz points out,⁵⁹ a single spike of a neuron is a discrete event, but information is encoded in the temporal patterns of spikes, not in individual spikes). Until these two facts are reconciled, it is hard to even begin to ask serious questions about the neural implementation of cognitive patterns.

8 Wrapping Up

To sum up, the question at the heart of the mind/brain is ultimately the question of how specific patterns of activity in specific cells or cell assemblies give rise to cognitive patterns, not the question of which brain regions are associated with cognitive functions. The mind is what the brain does, not where it does it. It helps, of course, to know what regions of the brain are especially active when the subject is performing a specific cognitive task, but at the end of the day we want to know what the special properties of these brain regions are; such properties should be characterized in computational terms, because that’s the only hypothesis we have about cognition (the computational theory of mind). On this view, as Poeppel⁶⁰ recognizes, the interpretation of localization is fundamentally recast: What must be localized is the circuitry that executes specific computations that cognitive scientists have motivated on independent grounds.

In his discussion of the language faculty, Andrea Moro⁶¹ has likened the sentence to a tapestry: what we observe (the surface string of words) is not what we should

be interested in. We should go behind the tapestry and figure out the intricate threading patterns that result in what we see. I find this an interesting analogy, and I would like to suggest that it should be applied to our brain images. We should go behind the patterns of activation and look at the patterns of connectivity in the brain's white matter; we should look at the chemistry of synaptic connections between neurons, and try to find the correlates of the atoms and molecules of our mental chemistry, and then possibly ask the question of whether the way the brain works explains why the mind works the way cognitive scientists identify. There is no guarantee we will ever achieve such a level of explanation, but before we even try, we should be clear about the sort of questions we should raise.

Note also that success in this enterprise requires first that we form robust links between knowledge of language and how it is put to use (the topic of Chapter 9); short of that, we will never be able to take full advantage of millisecond accuracy data of the sort brain imaging techniques make available.

In sum, for all the usefulness of the classic model and its main characters, Broca's area, Wernicke's area, and the left hemisphere, there is a desperate need for a new functional anatomy of language, one that is much more refined in cognitive/computational terms.

Homo Combinans

1 Reconsidering Basic Assumptions

Until recently, and, in many quarters, perhaps still to this day, whenever the topic of evolution of language was discussed, three things would happen: First, someone (say, the author of a paper on the topic) would bring up the ban imposed by the Linguistic Society of Paris in 1866¹ on all talk revolving around the evolution of language (on grounds that the issue invariably led to wild speculation). Second, once the author had recognized that the ban had been appropriate in those days, he/she would indulge in speculative reconstructions of our linguistic past (and I don't mean here reconstruction of ancient languages, as practiced by philologists, I mean something cognitively much more interesting: the emergence of the language faculty in our species), and would point out that since we know more about both evolutionary processes and the character of the language faculty today, we do not run the risk of repeating the sort of just-so stories that led to the 1866 ban (though many of these newer accounts, it should be said, turned out to be variants of the old stories). And third, since the comparative method plays such an important role in evolutionary studies, a discussion of the potential linguistic abilities of other species (our closest relatives in particular), especially their ability to learn a human language, would always be in order.

Far be it from me to deny the usefulness and importance of evolutionary studies focusing on mental faculties, language included. If the perspective from which we study such faculties is biological, as it has been since the cognitive revolution, evolutionary considerations cannot be ignored. As Tinbergen² and Chomsky made clear in their sets of foundational questions reproduced at the beginning of this book, it is reasonable to ask how mental faculties, once identified and properly characterized, emerged in the species. I also take it that the only game in town here is Darwinian (naturalist)³ in nature; creationism has no place if the goal is to render the object of study intelligible.

The problem with evolutionary studies on language (and other cognitive domains) is that the question of origin/emergence is almost invariably replaced by the question of adaptation and selection:⁴ What was language (or any other cognitive faculty) for? What led to its selection? What was its advantage? This narrower formulation of the evolutionary question may not only lead one to work with a smaller space of possibilities than evolutionary theory allows (as Darwin pointed out, “natural selection has been the main, but not exclusive means of modification”⁵), it may also not be the most productive question when it comes to cognitive capacities like language. Questions of adaptation turn out to be very subtle, even in the case of simple organisms and functions, let alone for complex capacities put to various uses like language, that are furthermore in some obvious sense unique to a species. This, combined with a very poor fossil record for cognitive abilities of extinct lineages, renders many of the key necessary facts empirically inaccessible for now and perhaps forever.⁶ Not surprisingly, adaptationist scenarios abound, and there is little in terms of testable hypotheses to tease them apart. All in all, although the evolutionary question is appropriate, the adaptationist question appears to a growing number of researchers, myself included, to be a very poor choice as a defining issue.⁷

As for the attempts to uncover linguistic abilities in other animals, they too have suffered from a poor selection of guiding questions. First, comparisons have been limited to the domain of communication. While this is quite understandable (it is hard enough to figure out the relation between language and thought in humans, and it is even harder in the case of animals who can’t answer questions directly), it is not at all clear that the language faculty is first and foremost a communication system.⁸ There is a long-standing alternative view, according to which the language faculty is primarily a system that enhances thought (with externalization being a side-effect, as it were); so comparing the human language faculty and communicative behavior in other species is a bit like comparing apples and oranges. No wonder they look very different. Second, most comparative studies approach the issue by framing the issue in very black-and-white terms, taking language to be something one either finds or does not find in other animals. Language, in other words, is taken to come as a package. Occasionally, “language” is decomposed into syntax, phonology, semantics, and the lexicon, but as we saw in the previous chapter, even this sort of decomposition of the language faculty is much too coarse. If it’s not the right level of description for human language, you can be sure it’s not the right level of comparison.

As a result of these two problematic assumptions, comparative studies have tended to be contrastive, and negative in their conclusions.⁹ They have shown that many species have evolved incredibly sophisticated communication systems, which are interesting in their own right¹⁰ but irreducibly different from human language.

Furthermore, it is now generally acknowledged that attempts to teach human language to members of other species have failed.¹¹ Although the popular science press is always eager to seize on stories of dogs learning words, or chimpanzees producing sentences, or parrots reproducing sentences, it is hard for me to understand why we even started running human-language-teaching experiments with other

species. If we take language to be part of the human biological endowment (as I hope the reader now thinks we should), there is no reason to expect that animals with a different biology from ours would be able to do what our distinctive biology makes possible for us to do. No one submits grants to teach human graduate students how to perform the complex waggle dance of bees, or how to catch bugs by emitting short pulses of high frequency like frogs, or how to echolocate like bats – for a good reason: everyone, correctly, assumes that this is beyond our biology (much like it is outside of our biology, but not that of other species, to perceive ultra-violet color). Why should human language be treated differently? Why are some scientists tempted to say that Ginger the dog learned 100 words like baby Roger did, while nobody is tempted to say that Olympic high-jumpers fly, like birds do?

2 Comparative Studies

I will not attempt to survey the various systems of communication that can be found in nature and that have been studied with great care by many researchers;¹² but if one were to try to summarize the major findings, I think it would boil down to this.¹³ Comparative studies have revealed that the communication systems of all other known animals are based on limited, fixed sets of messages, all of which are limited to the here and now. Typically, the set of possible messages cannot be expanded by combining elements to form new and different complex messages. Such messages also fail to display the sort of flexibility of use one finds in human languages.

For example, vervet monkeys¹⁴ have developed calls to warn against predators (there is a distinct call for the eagle, the snake, and the leopard, and one or two others, including one, I'm told,¹⁵ for humans – but the list is very small), but they are only used in the presence of a predator. Unlike words in natural languages, calls are clearly referential, and involuntary. Their communication system is thus much more reflex-like than ours (although they are not simple stimulus-response behaviors of the sort Skinner and his colleagues would have liked).

Bees¹⁶ have a complex dancing pattern that enables them to communicate information about food: how much of it there is, and how far it is from the hive, and how to get to it (using the sun as a navigation landmark). Computationally speaking, this is far from trivial, but notice that there is little creativity in the system. The system can't be recruited to "talk" about something else, and it can't be used to compare today's food source with last week's. It is also interesting to note that the units of the message being conveyed are not discrete (unlike words in human languages): information varies along a continuum, showing more or less food, but never exactly this much food; in this regard, the units of the bee dance resemble our approximate number system (cf. Chapter 8).

It is also interesting to note that in most examples of communication systems I know of, the system is strongly innate; it emerges without much relevant experience, and in some cases is available right at birth. Unlike human language, where

sound/sign–meaning pairs must be learned, animal communication systems aren’t learned in any interesting sense. As various researchers have noted,¹⁷ they may require a bit of fine-tuning concerning the precise condition of use of some signal (as appears to be true for monkey calls), or concerning “dialectal” adjustments (there are some dialects in bee communication), but young members of the species can’t acquire genuinely new signals, or change the basic repertoire in fundamental ways, unlike human infants.

Perhaps the only exception to this generalization is found in the context of songbirds.¹⁸ As has been known for a long time by ethologists,¹⁹ some birds develop their characteristic singing patterns spontaneously, with little or no relevant exposure, but other birds learn their songs from conspecifics like humans do. Some birds have a very limited song repertoire (sometimes limited to just one song), but other birds’ repertoires are much richer. Interestingly, birds learning their songs go through developmental stages reminiscent of what is found with humans:²⁰ they first go through a subsong stage that is very close to what we call babbling, then they progressively learn songs of greater complexity until they reach the adult stage. Song learning, as is so often the case in learning matters, is also subject to a critical period. In some species, songs have an intricate internal structure (with various recognizable chunks within the song, and “syllable”-repetitions, and so on), although the meaning appears to be always the same: singing is used to mark and defend one’s territory and attract mates (as is true of many communication systems, including vocal ones, all the way to the fish).²¹ No “new” meanings are conveyed by new songs (for birds with a rich repertoire), it’s always a variation on the same theme. Nevertheless, birdsong is probably the domain where comparison with human language has established the most interesting similarities (over and above obvious differences). Some²² have even suggested that songbirds may offer us a unique window into the neural basis of (some aspects of) language learning, on the basis of the fact that the song system appears to enjoy a left-hemisphere dominance, and a system shared by production and comprehension.

Still, across systems of communication, human language stands out as special in at least two respects, both of which contribute to the vast expressive power of human language. First, humans are able to memorize many thousands of words, each of which encodes a distinct piece of meaning using arbitrary sound or gesture. Second, humans are able to combine words to form sentences, making it possible to communicate infinitely many different messages, and providing the basis for human linguistic and conceptual creativity (cf. Chapter 8).

3 The Irreducible Difference

Although I have no doubt that other species can pair a sound (or other signal) and a concept – some may even be able to pair a human sound/signal and a concept (after all, some animals, such as dogs, have evolved to be sensitive to human

communicative intentions), we should not conclude from this that they have acquired words. Words in human language are not just sound/sign–meaning pairs; they have what Chomsky has recently called an “edge” property,²³ an ability to combine and recombine with other words to form an infinite range of expressions. The key difference is thus combinatorial in nature, down to the level of the word.²⁴ Furthermore, this key difference is responsible for a conceptual difference. As I remarked above, the meanings conveyed by other species are on the one hand confined to the here and now, and on the other, they seem to be confined to highly specific domains (food, predators, etc.). The latter point is very reminiscent of what we discussed in Chapter 8 in the context of core knowledge systems. There I followed a growing number of specialists in claiming that the mental life of animals is undoubtedly richer than we are often tempted to attribute to them. Their computational abilities are quite impressive, and other species are known to outperform humans at a variety of tasks. But even their strongest advocates will confess that the most striking thing about other intelligent creatures is their “narrow-mindedness,” as it were – their inability to step outside the bounds of the domain in which they excel and apply their intelligence to another, new domain, or combine two previously unrelated conceptual spheres. That’s where humans stand out. Animals are specialists; we are generalists. Humans can easily transcend the bounds of their sensory inputs, and form concepts that combine information from distinct modules. The net result of this ability is akin to what happens when hydrogen and oxygen atoms combine to yield something with such dramatically different properties from its constituent elements as water; the product is more than just the sum of its parts.

It is, I think, a safe bet to assume that human language lies at the heart of this conceptual prowess (as many have conjectured; cf. Chapter 8). It’s as if by turning a concept into a word, humans are able to detach this concept from its natural domain (module) and place it in a “neutral” workspace where other concepts, from distinct modules, can meet. In virtue of all being words, concepts of different shapes, with very different properties, come to share the same format, and can combine without restrictions. Humans have also evolved the capacity to make these new concepts known to others, via speech or sign.

All in all, these differences, however dramatic, are few in number. Computationally speaking, they boil down to: lexicalization, combination, and externalization. This minimal characterization of the source of the difference between us and other species pertains not just to systems of communication, but also systems of thought, and fits well with the general outlook (research program) articulated in an already famous paper by Marc Hauser, Noam Chomsky, and Tecumseh Fitch.²⁵ They were led to collaborate because each of them was convinced of the fruitfulness of the biological approach to mental capacities (as their writings attest),²⁶ and each of them was interested in addressing evolutionary issues, but they all agreed that debates surrounding issues of adaptation were more conducive to speculation than to empirical research. They asked themselves and each other how to move forward in exploring evolutionary issues, and came up with a collective answer that has since

then defined what strikes me as the best (most feasible) research program in this area, and it is this program that I want to outline in the remainder of this chapter.

4 A More Fruitful Approach

Hauser, Chomsky, and Fitch began their original paper²⁷ by stating that the question of the evolution of the language faculty must strike a proper balance between the fact that the language faculty is unique to our species and the fact that undoubtedly humans and animals share a diverse array of important computational and perceptual resources, many of which are likely to enter into the linguistic domain. Accordingly, the research focus should not be placed on “language” (or even on the “language faculty”), but instead, it should be placed on the computations underlying the system, and on the representations over which these computations are carried out. When we engage in comparative studies, we should be asking whether these specific computations and representations that serve human language are shared among species or unique to us.

They went on to make a conceptual distinction between the faculty of language in the broad sense and the faculty of language in the narrow sense. The faculty of language in the narrow sense is defined as that which is unique to language and unique to humans (it corresponds fairly closely to the notion of “competence” discussed in earlier chapters). The faculty of language in the broad sense includes the faculty of language in the narrow sense but also all the other resources that enable this capacity to be put to use; it thus includes all the systems that the competence system interfaces with, like the sensori-motor systems ultimately responsible for the externalization of language, and the conceptual systems that language connects to.

Hauser, Chomsky and Fitch stress that whether one studies language in the broad sense or the narrow sense, the focus should be on the computations involved; one should avoid talk of syntax, semantics, etc. – the coarse distinctions that render both evolutionary and neurological studies sterile. Each domain (syntax, semantics, etc.) should be decomposed into component mechanisms, which can then serve as a basis for investigation. For this reason, Hauser, Chomsky and Fitch warn that the concepts to be dealt with are unlikely to map onto traditional boundaries. That is to say, the domains of investigation will appear even less intuitive than what one may find in introductory texts, but as we already saw in the context of neuroscience in the previous chapters, this is the price to pay to make progress, for intuitive, or even traditional, scientific distinctions rarely cut nature at her joints (the history of science shows how often this has turned out to be true). Just like it is too vague to ask where syntax is in the brain, it is too vague to ask how syntax emerged in the species.

They also point out that, as a starting point, it may be best to focus on a characterization of the most basic and essential aspects of the language faculty – properties that everyone would agree on, as opposed to the sort of detailed aspects

that are at the center of theoretical debates. Furthermore, a premium should be placed on characterizing the core properties of language in a very general (jargon-free) way, so as to facilitate interdisciplinary dialog, and also comparative studies. For ultimately, the empirical challenge is to determine what was inherited unchanged from our ancestors, what has been subjected to minor modifications, and what (if anything) is qualitatively new. The point of the Hauser, Chomsky, and Fitch proposal is to articulate a series of questions that make comparative research possible and potentially revealing (a radical departure from previous works concerned with teaching human language to other species). The hypothesized computations should be formulated in such a way that their existence can be tested in experiments with non-linguistic species. This is undoubtedly the most difficult challenge that followers of the Hauser, Chomsky, and Fitch proposal will have to meet; on the one hand we want to avoid claiming that some particular aspect of language is specific just because it is characterized in an idiom that only linguists find adequate, but on the other, formulating properties that are specific to language in too general a fashion may give the impression that they are not specific at all.

To give the reader a sense of this tension, let me discuss the results of two recent experiments that take as a starting point the computational difference between finite-state machines and rewrite rules discussed in Chomsky's *Syntactic Structures* and reviewed in Chapter 4.

5 Interpretive Difficulties

Recently, Timothy Gentner and his collaborators²⁸ reported that they had been able to teach European starlings to recognize a difference between song patterns that seems to mimic what only phrase-structure rewrite-rule systems (as opposed to simple finite-state automata) can achieve. Gentner et al. taught the starlings a series of songs in which a number of “rattle” motifs were followed by a matching number of “warble” motifs, and other songs where the number of “rattles” was different from the number of “warbles.” After a huge number of training sessions (10,000–50,000 trials!) most of their birds could learn to distinguish songs of the first type from those of the second, even in songs they had not previously been exposed to. In formal terms, songs of the first kind had an A^nB^n pattern, whereas songs of the second kind were A^mB^n (where the number of As and Bs were four or less), exactly the sorts of patterns that lie beyond the computational limits of finite-state machines. But should we conclude from this that European starlings have learned the distinction by making use of the computational system that we think gives human language its distinctive structural character (nested dependencies)? Unfortunately (or perhaps, interestingly), matters are more complex. As many researchers have pointed out,²⁹ the birds may be using a completely different (and far less powerful) system to track down the number of As and Bs. They may simply count the As, count the Bs, and compare the sets. This could be done by relying

on a finite-state machine supplied with a counter. Since we know that birds, like virtually all other animals ever tested, are capable of recognizing the cardinality of small groups (subitizing; cf. Chapter 8), this possibility is very real. How to tease the two possibilities apart, without making implausible assumptions about the performance systems (memory limits) that must be taken into consideration in any experimental design, has proven (so far) impossible.

The Gentner et al. experiment received a lot of attention in the literature because it had been inspired by a very similar experimental design that Fitch and Hauser³⁰ tried on cotton-top tamarins and human babies, with a different result. Fitch and Hauser reported that the cotton-top tamarins failed, but the human babies succeeded in learning the pattern (A^nB^n) that goes beyond the computational resources of the finite-state machine.³¹ Although the result of Fitch and Hauser's experiment could be taken to indicate that our closest relatives indeed lack the computational resources that form the core of our language faculty, as many suspect, we must make sure that the human babies didn't resort to the counting trick that starlings may employ, or some other alternative route.

In sum, although the Hauser, Chomsky and Fitch perspective leads to a more empirically minded (hence productive) research program in the domain of evolutionary/comparative studies, there remain inherent difficulties in carrying out the relevant experiments across species; difficulties that have proven very hard to overcome so far.

Perhaps for this reason, there isn't yet a leading hypothesis to report on in a book like the present one, but the fragmented view of the language faculty promoted by Hauser, Chomsky, and Fitch naturally leads to the expectation that our unique ability called language will be a mixture of old computational parts and new computational parts. And perhaps more interestingly, even the new parts may be the result of some minor reconfigurations of pre-existing computational capacities, or some capacity that existed in just one small corner of the computational mind and then got generalized. All these possibilities are open, and they fall within the range of evolution conceived of as "descent with modification," Darwin's favorite characterization of the evolutionary process (in fact, he much preferred the term "descent with modification" to "evolution"), as it stressed both continuity and divergence.³² It is from this perspective that one should re-examine the process by which concepts get lexicalized, and the process by which words, once formed, can be combined and recombined freely.

We should also examine the possibility, very real from the perspective outlined here, that once the key novelties evolved, the older parts with which they interfaced were in turn modified: in Chapter 8 I already suggested (based on works by Elizabeth Spelke and others) how a mind equipped with language (now meaning a mind equipped with the computational resources to uproot concepts from their modular bounds and combine them at will) can acquire a new number sense that goes beyond what is available to other animals; the same could be true of other seemingly unique abilities such as music, mind-reading (theory of mind), etc. Old, primitive capacities may receive a boost from the emergence of this new range of

computations. The whole mind may have been reconfigured as a result of a few small changes.³³ Likewise, we should expect that if language draws on older computational tools, it should inherit some of their inherent signature limits.³⁴ While there is no doubt that the emergence of language has expanded our thought systems, one should not lose sight of the fact that their foundations (like everything else about our body) remain profoundly entrenched in animal cognition (and anatomy).

While addressing these questions of old and new, we should always bear in mind that some of the old parts may be very old indeed. Some computational capacities may have been recruited or re-activated in our lineage a long time after they were developed and subsequently silenced. It is quite plausible, in my view, that the computational tools necessary for the externalization of our language faculty may lead us back to songbirds, as Massimo Piattelli-Palmarini and Juan Uriagereka have suggested.³⁵

Consider vocal imitation, the ability to flexibly recreate novel utterances after hearing them produced by another individual. This is a well-developed capacity in human children and it is a crucial prerequisite for spoken language. Without it, our giant lexicon would remain the stuff of dreams. Now, although well-developed vocal imitation is observed in many nonhuman species, including birds, cetaceans (whales and dolphins), seals, and perhaps bats, it has not been observed in any non-human primate, including chimps, despite intensive investigation.³⁶ While some birds can imitate arbitrary sounds, including human speech, chimpanzees fail to do so even with intensive training. Could it be, then, that some capacities crucial to the evolution of our language faculty got revived, as it were, from neural substrates we share with birds?

Our language faculty may then be a mosaic of very old parts, a few recent innovations (themselves maybe mere modifications), and basic chimp-like conceptual resources; it's the combination of these that resulted in this unique cocktail we call *Homo sapiens*, perhaps more aptly characterized as *Homo combinans*,³⁷ to emphasize our remarkable combinatorial capacity. It now becomes interesting again to see if animals may be able to learn some aspects of language, though now crucially characterized in computational terms, and with the possibility that the computations may be present in another cognitive domain.

6 Coda

I hope it is clear to the reader that the research questions formulated by Hauser, Chomsky, and Fitch enable us to move from unproductive debates about adaptationist scenarios to more collaborative empirically focused discussions, where comparative data are thrown into a much more positive light.

Of course human language is adaptive. By allowing us to form and communicate an endless variety of thoughts, it's clearly an adaptive trait. We are such a thoroughly linguistic species that it is hard to see how we could have remained alive as a species

if language somehow didn't confer an advantage. But the fractionation of the language faculty that Hauser, Chomsky, and Fitch ask us to consider recasts the adaptationist question. It moves the issue from human language as a whole and the ways it is put to use to the particular components that make it possible, to the computations involved. Were these adapted specifically for language, or were they adapted for something else and later got recruited for language, or were they simply a side-effect of some unrelated change, a bit like the visual illusions that necessarily form given the right environment? To answer these questions we would need to know much more about the history of our species, and about the various physical constraints and accidents that channeled our evolutionary process. Our upward posture, changing diet, and many other factors no doubt contributed to our distinctive cognitive profile. But perhaps we will never be able to reconstruct the right chain of events, as Richard Lewontin never tires of pointing out.³⁸ The usefulness of evolutionary questions may not lie in trying to reconstruct the past that is forever lost; rather, the greatest value of an evolutionary perspective,³⁹ with the comparative method at its core, may be to provide a theoretical framework within which to frame modern empirical research, to act as a guide to seek new data and generate testable predictions, and to refine our understanding of the computations involved in the domain of language and elsewhere. Rather than answering how our language faculty really emerged, it may be one of the best ways to find out what our language faculty really consists of.

12

Computational Organology

Given the range of phenomena reviewed in this book, I think it's fair to say that the human language capacity (and other cognitive faculties) must derive, at least in part, from properties of the genome. In this final chapter I would like to address issues touching on genes and behavior, taking language as a case study once again. The discussion will allow me to mention interesting disorders and cases of so-called dissociation; it will also allow me to revisit the issue of "specificity" and stress the need for a computational approach if one is to relate mind and brain, behavior and genes. The latter point will directly relate to some of the conclusions reached in the previous two chapters.

1 Linguistics and Genetics

Let me start with genes. Although it has long been suspected that our genetic make-up is somehow ultimately responsible for our cognitive capacities, it was only recently that concrete examples of genes involved could be identified, most dramatically in the case of disorders (as is the case with neuroscience, it's when things break down that we can try to infer how things work in the normal case). Today, we know a great deal in some instances about which genes are associated with which specific disorders, but we have little idea about why those genes have the specific consequences that they do, for language and so many other traits. The gulf that exists between our understanding of the genetic causes and the behavioral outcomes of developmental disorders is immense.

One thing we can be sure of is that the idea that individual genes control individual behavioral traits is wrong (even if the media still use phrases like "*the* gene for *x*"). Genes play a far more complex role in regulating the actions of other genes, and the synthesis of proteins, which, among many other things, give rise to specific patterns

of connectivity among neurons, resulting in the specific human brain, somehow equipped with cognitive faculties like language. But the road from what genes do to what the brain does is tortuous and far from direct, or obvious. Besides, there are simply not enough genes available for each gene to control an individual trait.¹

As Colin Phillips² has pointed out, this state of affairs makes the observation that specific genetic disruptions lead to specific cognitive disorders all the more puzzling. As he notes, if it is true that genetic disruptions cause children to have special difficulty in, say, inflecting verbs properly for tense, and if we can be fairly confident that no gene codes for tense inflection, it's all the more puzzling that specific areas of language turn out to be more vulnerable than others. Even more puzzling is the fact that different genetic disruptions appear to lead to often similar areas of vulnerability. Thus, in the case of language, inflection and structures involving non-canonical word order (passives or objects "preceding" subjects in relative clauses in English) appear to be particularly vulnerable across a wide range of disorders.

We don't know why any of this should hold. We don't even have good intuitions about what might be going on, I think. These are questions that will likely have to wait a long time before we can even get our minds around them. But as I will try to make clear in this chapter, standard practice in modern cognitive science is sure to have an interesting role to play, for it ought to go without saying that in order to understand how specific neural structures, and the genes giving rise to them, support cognition, we need to have a proper understanding of what cognition is, and our very best bet is the computational theory of mind, the main character of this book.³

2 A Case Study: FOXP2

Perhaps the best illustration of all the issues just raised comes from what some have called the saga of the FOXP2 gene.⁴

About 20 years ago an entire English family, known as the KE family, was brought to the attention of neurologists and speech therapists in London. Over three generations, roughly one half of the family members showed a remarkable language deficit, which, as Massimo Piattelli-Palmarini and Juan Uriagereka⁵ point out, presented experts with the textbook characteristics of an inheritable condition – the sort of example people would use to illustrate the logic of Mendelian genetics to students. Not surprisingly, the KE family case made it into textbooks within a few years.⁶ But matters are, of course, very complex, at all levels (mind, brain, behavior, and genes, including evolution of the genes).

The affected members of the KE family were examined by a variety of specialists.⁷ As far as anyone can tell,⁸ they are sociable; they are aware of their deficits, and go to great lengths to make themselves understood. The level of general intelligence of family members is roughly average (although this remains a controversial issue), and it is generally agreed upon that there is a double dissociation in the affected members of the KE family between general intelligence and their linguistic

impairment: that is to say, among the affected population, there are cases of normal to high non-verbal IQ with heavy linguistic impairment, and of low non-verbal IQ with only minor linguistic impairment.⁹

In order to evaluate their deficits, various tests were administered (some of these tests are also used with aphasic patients). Affected members of the KE family appear to have great difficulty in repeating “non-words” (non-existent, but plausible words, like *hampent*), and understanding complex syntactic and related morphological nuances, such as constructions involving relative clauses, subject–verb agreement, and plural marking on nouns.

Exactly how to characterize this deficit has been a hotly debated issue ever since linguist Myrna Gopnik’s 1990 original hypothesis in terms of “feature blindness” (inability to deal with some features/properties of language like verbal and nominal inflection), with some arguing for a case of Specific Language Impairment (SLI) (a family of developmental language disorders, affecting both expressive and perceptual aspects of the language faculty),¹⁰ others¹¹ claiming that the central deficit is not specifically linguistic, but rather a severe “oro-facial dyspraxia” (articulatory deficit) combined with developmental cognitive deficits which, somehow, affect the learning of “fine sequences of subtle motor controls.”

Despite the controversy as to what exactly the syndrome boils down to, various specialists tried to track the genetic and brain correlates of the deficit, and succeeded.¹² In 1997, region 31 of chromosome 7 was singled out as the locus of the deficit, and was named SPCH1. SPCH1 turned out to be a member of the FOX (forkhead/winged-helix replicator gene) family, of which several other genes are known all across the animal world. SPCH1 was then relabeled FOXP2, as it is known today.

In 2001,¹³ the exact locus and nature of the mutation affecting the KE family was determined. In one copy of the gene only (the maternal one), one finds a mutation that is not found in non-affected individuals in the KE family, nor in hundreds of controls for the normal population. This mutation alone turns out to be necessary and sufficient for whatever deficit the affected KE family members ultimately exhibit. It has been said¹⁴ that the result of the genetic studies on FOXP2 is as detailed a result as one can hope for in figuring out the contribution of genetics to the analysis of a phenotype; it is, in the words of Piattelli-Palmarini and Uriagereka,¹⁵ “a geneticist’s dream”.

It should be said at this point that the activity of FOXP2 is quite complex.¹⁶ Technically, the gene codes for a transcription factor, a regulatory protein that modulates the transcription of other genes from DNA into RNA. In less technical terms, the FOXP2 gene is like a big boss; it controls the activity of other genes which are ultimately related to a variety of organs, including the gut, the heart, the lungs, and, yes, the brain. (I should point out that the deficit found in the KE family does not affect, as far as one can tell, the development of the various other organs in which FOXP2 is implicated.) Interestingly, and somewhat counterintuitively, greater activity of the protein(s) produced by this gene means less activity in the genes that they regulate. In other words, a defect in this gene causes more of other proteins to be synthesized, and/or a greater quantity of those proteins.

This very brief overview of what we know about the gene should already be enough to dispel any thought of FOXP2 being “the language gene.” FOXP2 is just one element of a complex pathway involving hundreds of other genes,¹⁷ and it’s too early to tell how special the role of FOXP2 is. But the problems exhibited by the affected members of the KE family, touching as they do on important components of linguistic computation (agreement, inflection, etc.), nevertheless suggest that we should take this gene very seriously.

Just like the situation at the behavioral and genetic levels, the situation at the neuronal level is anything but clear. Some brain areas have been found to correlate with the behavioral deficits, but, as the reader familiar with Chapter 10 may expect, each of the brain areas identified has been associated with multiple functions, and it strikes me as futile to even try to summarize what is surely too speculative.¹⁸ At this point all we have are interesting brain–behavior correlations, nothing more.

In recent years, FOXP2 has also figured prominently in evolutionary studies related to language, following a 2002 study,¹⁹ where the phylogenesis of *Foxp2* (conventionally, only the human version of the gene is all capitalized) was reconstructed across the mammalian world. *Foxp2* turns out to be an old, remarkably conserved gene. Only 3 (out of 715 possible) point mutations exist between the version of the gene in humans and the one in, say, the mouse (distant about 150 million years, having had a common ancestor some 75 million years before the present). However, the result that has fueled a lot of speculation is that two of those point mutations exist between humans and chimps (distant some 11 million years). In other words, two recent mutations in the expressed region of the gene took place in our lineage, while things remained pretty much stable in our joint primate lineage with mice for the best part of our common evolution.

As one can imagine, this result was seized by some to argue for specific evolutionary scenarios concerning language. It was suggested²⁰ that the effect of the gene (by hypothesis, language) was heavily selected for, hence the two mutations in what is otherwise a remarkably conserved gene. But the complexity of the genetic role of FOXP2 suggests that we should tread very carefully here. The two mutations and their dates may well be facts, but they are subject to multiple interpretations. Since FOXP2 plays many roles (to date: in the spleen, muscle, liver, skin, uterus, eye, mammary glands, testes, kidneys, lungs and intestine – at the very least, and quite apart (or so it seems) from its clear expression both in the brain and the “speech organs” at large (ear, tongue and trachea)), any of these roles may be related to the two mutations, and the linguistic effects may be mere side-effects. At the present we just don’t know.²¹

In the end, this is the overall conclusion of the FOXP2 saga for now: nobody knows what’s really happening, deep down, with the affected members of the KE family, at any level (mind, brain, behavior, or genes), but this, I hasten to add, is a common situation when “things get interesting.” Everyone should realize that, however complex the current state of affairs may be, we have here the best instance to date of a possible bridge between language and genes, something that Lenneberg was perhaps dreaming about, one that, for my purposes, serves as the very best illustration of all the issues and challenges involved in integrating mind and brain.

Piattelli-Palmarini and Uriagereka²² are right to stress in their review that among the many things that need to happen in the context of FOXP2, or any other such case to be discovered, one thing is clearly required: linguists and cognitive scientists should offer their expertise in seriously analyzing the available data and even building tests to acquire new information. Cognitive science won't be sufficient, but it will be necessary. It should also be clear that in order to be useful, results in cognitive science should be formulated in computational terms.

3 Research Readjustments Needed

This was already the conclusion reached by David Poeppel,²³ discussed in Chapter 10, and it is also at the heart of Hauser, Chomsky, and Fitch's decision²⁴ to break down the language faculty into its component mechanisms (Chapter 11).

We should abandon once and for all the idea that “syntax” (or any other corresponding notion) is in the brain. As Poeppel notes,²⁵ if anything is localized in nervous tissue, at least at the level of cellular ensembles or columns, it will be *elementary* computational functions. Poeppel goes on to say:

Linguists and psycholinguists owe a decomposition (or fractionation) of the particular linguistic domain in question (e.g. syntax) into *formal operations that are, ideally, elemental and generic*. The types of computations one might entertain, for example, include concatenation, comparison, or recursion. Generic formal operations at this level of abstraction can form the basis for more complex linguistic representation and computation.

Accordingly, cognitive scientists, and linguists in particular, should develop the set of elementary representations and computations.²⁶ The challenge to their neurobiological colleagues will then be clear and precise.²⁷ Specifically, they will have to ask what types of neuronal circuit form the basis for the hypothesized computations in question.

As I already pointed out in Chapter 10, such elementary computations may or may not be localized to one place in the brain; they may be distributed over different brain areas, and may even be instantiated multiple times in different places. In this regard, the model for linguists should be vision science.²⁸ Although much remains to be understood there too, especially when it comes to what is known as higher-level vision (3D-shape computations, etc.), we already know that our vision faculty is a collection of small computational components – indeed, the vision faculty is fragmented in ways that fail to match our intuitions. Thus, what is called “early vision” (or early visual information processing) comprises at least three stages, or levels/circuits. The first is the primary sensory level (including the eye and the primary visual cortex at the back of our brain). Second is the primary unimodal level, of which there are two for vision: one circuit that processes information concerning where an object is (depth, motion, position), and the other focusing on what that object is. Once that information is computed, it's integrated in a third

level, the unimodal association level, concerned with a further stage of information integration (color, motion, form of objects). Finally, all that information has to be integrated with information from other modalities in what are called multi-modal associations areas. Without going into any detail, it's already clear that when we process a visual scene, different aspects of objects are processed in different parts of the brain: the initial processing takes place in an area of the brain called V1, stereo vision takes place in V2, distance is computed in V3, color in V4, motion in V5, and object position in V6. All of that information ultimately gets integrated, but each aspect of what we perceive appears to have its own little computational module. We should expect nothing less in the context of "language" (or even "syntax," "phonology," etc.).

Going back to FOXP2, as with brain loci, it is doubtful that genes map neatly onto intuitive aspects of language. More plausibly, different brain areas, controlled by a hierarchy of genes, are likely to offer the means to perform the sort of elementary computations that cognitive scientists will recognize as the true ingredients of cognitive faculties like language. Assuming that this is indeed the case (recall this is just a hypothesis), we should be ready for the (very real) possibility that sometimes the most interesting counterparts of each one of the basic component of our cognitive faculties may be found in very different species, belonging to genera that are quite distant from ours. As Piattelli-Palmarini and Uriagereka²⁹ remind us, it is useful to bear in mind that the essentials of genetic recombination have been best evidenced at first in the study of bacteria and the fruit-fly, the mechanisms of embryonic development in the study of the sea urchin, and the dynamics of the nervous impulse in the study of the squid's giant axon. And, they also point out, in recent years it has become standard to explore mental retardation in children by studying specific mutants in drosophila. Obviously, we humans did not directly evolve from any of these species, but there are enough commonalities in particular organs and functions to make these model-systems interesting analogs, and thus looking into what a gene like *Foxp2* may be doing in other species may shed invaluable light on FOXP2.

For this reason, work on mice and birds has received a lot of attention recently.³⁰ Since we can manipulate them with advanced gene knock-out technology (whereby certain genes are made inoperative), mice are the first good candidate for the exploration of this issue. Interestingly, the relevant allele of *Foxp2* in mice appears to be involved in how pups ultrasonically communicate with their mothers.³¹ Needless to say, this is nothing like the way we communicate via language (and it is already telling us that whatever is at work, it is likely to be computationally more abstract than we think), but it may be that mice will tell us more about some aspect of the language faculty than our closest ancestors (for which a quasi-linguistic role for the gene has yet to be found) would. It may well be that the brain circuits that *Foxp2*/FOXP2 regulates in mice and humans are simply not active in, say, the chimpanzee. "Surprises" of this sort abound in the molecular biology literature, and cognitive scientists should be ready for them.

Experiments on songbirds may be highly revealing too. A number of songbird species, like zebra finches, are relatively easy to experiment with in the lab, and, as we saw in the previous chapter, there are interesting developmental and possibly neural

similarities between songbirds and us. Brain studies have long shown how the bird brain presents two relevant circuits, one for song acquisition, and a differentiated one for its performance. Interestingly, *Foxp2* is expressed in both, especially in a region called Area X that connects the acquisition and production circuits.³² It is expressed in the acquisition circuit while the male bird is acquiring the song, and later on in the production circuit as he sings it. Again, as was the case with the mouse, birdsongs are quite different from human languages. For one thing, they do not have the sort of rich, structure-based semantic import human language has. Whereas we get different meanings by recombining words, the variations in birdsongs have little meaning attached to them, so far as we know. But as Piattelli-Palmarini and Uriagereka³³ have suggested, human language and birdsong may share abstract computational properties. Perhaps the most obvious one is the fact that just like our hierarchical mental structures have to be linearized, the complex song patterns that ethologists have revealed must also be squeezed out into a one-dimensional sound.

Could it be, as Piattelli-Palmarini and Uriagereka ask, that our species has recruited an ancient gene with a relatively ancient function (one of many) to, in their words, “help us squeeze our thoughts out into the air-waves”?

Right or wrong, Piattelli-Palmarini and Uriagereka’s hypothesis has the advantage of focusing on computational properties, and I agree with them that linguists (and other cognitive scientists) have to look in this direction. They should be prepared to examine how their findings regarding mental structures could be integrated with results in comparative genomics, and they should be ready for the possibility that molecular mechanisms may challenge their theoretical analyses. As Piattelli-Palmarini and Uriagereka say, in the end, the issue is to be ready for what amounts to a new step in the sciences of language and cognition more generally, with concrete hypotheses that can be evaluated and improved on, as new data emerge. Needless to say, the very same message is valid for geneticists, neurophysiologists, speech pathologists, and evolutionary biologists, who should be aware of the discoveries made in the past 50 years in cognitive science.

Piattelli-Palmarini and Uriagereka further note that different sciences develop along different trajectories, and it is perfectly normal that when they actually meet for the first time, even if they have been expected to meet for a long time, results appear to be incommensurable. But, as we saw in this part of the book, there is mounting evidence to suggest that the best way to overcome this incommensurability and unite mind and brain will involve paying more attention to computational aspects, and a willingness to decompose standard cognitive computations into more elementary and general parts.

4 Specificity Revisited

Inevitably, this will lead to a reconsideration of the issue of specificity. As I pointed out at various points in previous parts of this book, cognitive science, since its inception in the seventeenth and eighteenth centuries, has been divided between

those who argue that the mind is a collection of special-purpose machines and those who favor domain-general processes. Although we saw many reasons to favor the special-purpose machine (modularity) view, the program that David Poeppel has labeled “computational organology,”³⁴ and that seems to offer us some hope toward unifying mind and brain, genes and behavior, suggests that a looser sense of specificity (and modularity) may be in order. If the goal is to identify the computational primitives that enter into cognitive faculties, it is likely to lead to the identification of “generic” processes (“combine,” “linearize,” “concatenate,” etc.), shared across many components. It is indeed likely that many circuits are recruited for many different cognitive functions and that it is the collection of several generic processes that gives rise to the specific faculties and modules identified by cognitive scientists.

From a Darwinian perspective, this should hardly be news. Evolution, as François Jacob pointed out, is a tinkerer.³⁵ What is new is very often a rearrangement of old parts in novel configurations. By far the most common source of novelty in biology appears to be the result of duplication and divergence,³⁶ supporting Darwin’s favorite characterization of the evolutionary process: descent with modification.³⁷ We therefore expect a lot of neural recycling,³⁸ with pre-existing brain circuits recruited for new activities.

From this perspective, it comes as no surprise that areas like Broca’s area cannot have a single cognitive function assigned to them. Specific computation is unlikely to map onto specific cognitive function (though dramatic modification of an ancestral computation may). The specificity of cognitive domains, for which there is massive evidence, will have to be reconstructed from shared (or duplicated) computations, the way they interact, and the way they interface with other systems. This is exactly what Hauser, Chomsky, and Fitch speculated in the domain of the language faculty: the faculty of language in the narrow sense (what is unique to human language) is likely to amount to perhaps a single computational novelty coupled with specific pathways to other cognitive systems. In the words of linguist Luigi Rizzi,³⁹ it may be

that much of the specificity of the language faculty is linked not so much to the use of special computational devices, but rather its place in the “topography” of the mind, to the overall function it must perform (relating sound and meaning), [and] to the contiguity with the other systems it must interact with at the interfaces.

I suspect that much of the work in the immediate future in cognitive science will be devoted to reconstructing the specificity of mental organs in terms of more basic computational and representational primitives.

By paying attention to this new possibility of characterizing modularity,⁴⁰ we could perhaps begin to make sense of what look currently like paradoxical situations, such as the one reported by Aniruddh Patel in the domain of music.⁴¹ Patel is an expert on music cognition, and the author of the most comprehensive survey to date of what we know about this other remarkable cognitive capacity of humans.⁴² As in the case of language and other mental organs, our music system is known to break

down in various pathological circumstances. Just like aphasia in the domain of language, amusia refers to situations where music is impaired, but virtually everything else about the mind is intact. Such cases of dissociation suggest that music is a specific system, related to, but distinct from other mental organs. At the same time, as Patel reports, brain imaging results show a high degree of overlap in brain region between music and language in normal subjects. How could this be? How could systems that overlap give rise to pathologies that can be completely dissociated (amusia without aphasia and vice versa)? Perhaps the way out of this paradox lies in a proper characterization of the computations involved and the patterns of connectivity, taking the notion of “descent with modification” seriously.⁴³

Computational organology, with its emphasis on primitive and generic computations, does not jeopardize results like those of Andrea Moro⁴⁴ and colleagues reported in Chapter 10; the fact that Broca’s area is active in highly specific tasks is consistent with the idea that some areas of the brain are dedicated to different types of computation, independently of particular content domains, although said computations may be more central to some cognitive domains than others.

Computational organology appears to be tailored to make sense of the high degree of specificity noted in some disorders, such as Heather van der Lely’s characterization of Grammatical Specific Language Impairment (G-SLI).⁴⁵ As already noted in the context of the KE family, Specific Language Impairment refers to a family of disorders which are likely to stem from different underlying causes, but all of which have some effect of linguistic behavior. Among the various kinds of SLI, van der Lely has focused on those instances that affect “grammar.” Although van der Lely still works with notions like grammar, syntax, and phonology, it is clear that the deficits target specific aspects of these broad domains (e.g., tense-marking, agreement, displacement, etc.), and one can be confident that a more fine-grained computational description of the various disorders would lead to better diagnoses, and thus better treatments.

5 Final Words

To sum up this chapter and this last part of the book as a whole, let me stress that the ultimate goal of the program outlined here, ultimately inspired by the success of the computational theory of mind, is to have theoretically precise, computationally explicit, biologically grounded explanatory models of cognitive capacities, be they unique to humans or not. To arrive at a seamless transition across levels of description (genes, brain, mind, behavior), questions about brain localization or evolutionary adaptations to specific cognitive niches will have to make way for a more computation-oriented approach. Since computations are the bread and butter of modern cognitive science, one can safely predict that the role of cognitive scientists will be a major one in the near future. Although much, much work remains to be done, there is already one positive lesson we can draw from the discussion in

this last part of the book: the key to progress lies in the decomposition of complex behaviors and cognitive systems into their component parts.

The findings of cognitive science, once formulated at an appropriate level of abstraction, will provide the boundary conditions that the neurology of the future must satisfy. Exactly how this unification/integration will take place, nobody, of course, really knows. Note that I use terms like unification and integration, not reduction.⁴⁶ Like physicists who posit theoretical entities like the neutrino, which later on get detected in experiments, cognitive scientists certainly hope that the mental structures they uncover will ultimately be verified at the neuronal level. Perhaps in some cases some psychological laws will be deduced from the substrates that implement them, as in the case of Weber's law, which we discussed in the context of our number sense. Stanislas Dehaene⁴⁷ has suggested that the law stems from the increasing width of the neuronal tuning curves as the numbers get larger. He has also suggested that the distance effect also discussed in the context of our number sense comes from the overlap between populations of neurons that code for nearby numbers. But we shouldn't expect that such reduction, if tenable in this case, will carry over to all the findings of cognitive science. Reduction of one domain of knowledge to another has been extremely rare in the history of science. Each level of description is legitimate, and each provides a degree of intelligibility that is valuable.

Consider the result of the experiment by Andrea Moro⁴⁸ and colleagues showing that the brain responds differently to UG-consistent rules and to UG-non-consistent rules. Such data from neuroscience are very valuable, but in and of themselves (as Moro and colleagues readily acknowledge) they are just as suggestive (not more) than the cross-linguistic studies that led to the notion of parameters and to specific hypotheses about what is and what isn't part of Universal Grammar (cf. Chapter 5). Likewise, the neuroscience evidence is not more compelling than the study that inspired Moro and colleagues. Two decades or so ago, Neil Smith and Ianthi-Maria Tsimpli⁴⁹ examined Christopher, who suffered from severe mental retardation but could learn languages much faster and much better than normal individuals – a linguistic savant, as such patients are known. Smith and Tsimpli attempted to teach a made-up language to Christopher where the rules were outside what linguists would consider the bounds of Universal Grammar, the UG-non-consistent rules in Moro et al.'s experiment. Although Christopher mastered 20 or so languages, he failed to learn the language concocted by Smith and Tsimpli. I hope the reader will agree with me that the result of the Smith and Tsimpli experiment is just as dramatic as the results of cross-linguistic investigations, and the result of the experiment by Moro and his colleagues. Not one result stands out as more compelling; in fact, taken as a whole, the experiments benefit from one another. For here lies the real strength of any scientific enterprise: corroborating evidence, coming from various domains. What I have shown in this part of the book is that by taking a computational approach, such corroborating evidence is more likely to be found, but, as Yogi Berra would have it, prediction is hard, especially about the future, and only time will tell if computational organology is as pregnant a field as I think it is.

Epilogue

Cognitive science is a field where questions still vastly outnumber answers. For this reason the discussion in the previous chapters has sought to provide the reader with a series of considerations that will enable her to grasp the essential character of the enterprise, as opposed to a report on specific results. As I stated at the outset, my goal was primarily to provide a structure for interpreting discoveries yet to be made. But this is not to say that nothing has been achieved over the past 50 years of intensive scientific investigation. My goal was also to do as much as I could to help the reader appreciate what we already know. There have been lots of exciting results, and these have led to new questions which, even if they lead us to hypotheses that move further and further away from our intuitions, are too attractive to ignore, as they touch on some of the deepest themes in intellectual history.

1 Stressing the Importance of Mental Structures

Cognitive science is first and foremost an attempt to study what one might call “cognitive mechanics.” According to the computational theory of mind, to think is to manipulate symbols in a particular manner, and a major task for cognitive scientists is to figure out the ways in which symbols can be manipulated so as to give rise to the various mental domains such as language, music, mathematics, vision, and so much more. Throughout this book I have used language as a case study because I think the “body of doctrine” concerning our language faculty is sufficiently rich, and the results already attained sufficiently robust, to allow us to touch on all the big questions in cognitive science.

Looking back at the preceding chapters, I feel that in many ways I have in some cases merely paraphrased, in some other cases amplified or refined, many of the ideas that were already discussed in chapter 1 of Noam Chomsky’s *Aspects of the*

Theory of Syntax, entitled “methodological preliminaries.”¹ Written over 40 years ago, this is a testament to its lasting influence.²

A few years ago Niles Eldredge³ tried to summarize Charles Darwin’s vision, and formulated it in terms of patterns. Darwin saw patterns that were crying out for explanation: he saw patterns of resemblance between modern species and fossils, and he saw patterns in the distribution of organisms (distinct forms of otherwise similar species, etc.). The same could be said about Noam Chomsky’s vision. Chomsky saw patterns that were crying out for explanation. He saw patterns found in all languages, he also saw patterns in the way languages differ, and, finally, he saw patterns in the way languages are acquired.

Chomsky understood, like the Rationalists had done more than two centuries before him, that an internalist (these days called “biological”) approach to the problem was the only way to illuminate the nature of the language faculty. He understood that the child’s brain is sensitive to only a certain range of linguistic structures in much the same way as the human eye is sensitive to only a restricted range of the electromagnetic spectrum (i.e., only to the continuum of colors between ultraviolet and infrared). Chomsky also recognized, like Descartes and Humboldt before him, that a central feature of human beings is their ability to make infinite, creative use of finite means, and that even if creativity itself may remain a mystery to us, at least the finite means making it possible could become the topic of serious inquiry in light of mathematical developments concerning the nature of computation in the first half of the twentieth century.

With Chomsky the focus of linguists moved away from sentences (and languages) and toward the minds of sentence (and language) users. The goal became one of identifying the innate computational proclivities, the biases that structure the signal in both language acquisition and language use. Building on the works of many experts (only a fraction of which I could include in the notes and bibliography), I have argued that this computational-mentalist focus is our best hope to integrate genetics, neurobiology, developmental studies, and linguistic/cognitive theory. We are still far from having achieved this unification, but at least we know of a few places where we can begin to look.

As linguist Ray Jackendoff has noted on various occasions,⁴ the deep concern with the fine details of linguistic structure is what distinguishes the investigation of language from other subdisciplines of cognitive science. Complex structure certainly exists in other cognitive domains. In vision, viewers structure the visual field into groups of objects in specific configurations, some in the background, some in the foreground (the figure–ground distinction that Gestalt psychologists used extensively), some in motion, some not, and so on. Similarly, what is known as episodic memory is supposed to encode particular events in one’s experience. As far as one can tell, such events must be structured (“represented”) in specific ways: spatially, temporally, etc. So it’s quite clear that the mind is structures all the way down (and mental rules behind them). However, Jackendoff is right to note that there is no robust tradition of studying the mental structures involved in vision and episodic memory, as there is in language. It may well be that this is one reason why the concerns of

linguistics often seem distant from the rest of cognitive science, but as I emphasized in the last part of this book, it's this focus on mental structure that is our best tool to probe the mind and the brain, so that's the focus we should advertise.

Fortunately, some have grasped the significance of this focus on mental structures and the rules giving rise to them, and have begun to elaborate parallel theories of other cognitive capacities (perhaps not surprisingly, this was the effect of close interaction with Noam Chomsky).

As a final illustration of the reach of the linguistic model, I would like to briefly discuss the recent approaches to our music sense and our moral sense. My hope is that the success of the "linguistic analogy" in these two domains will motivate researchers in other domains of cognition to adopt the same set of guiding questions.

2 Extending the Linguistic Model

Recently, a number of researchers⁵ have argued that at a certain level of abstraction there are striking parallels between the exercise and development of linguistic competence and the exercise and development of moral competence ("moral sense"). Thus, a striking fact about humans is that they demonstrate a quite sophisticated moral sense from a very young age. Very young children have the capacity to recognize moral rules and to distinguish these from conventional rules. They recognize that moral rules are independent of any particular authority and are associated with notions like harm or injury. Young children also understand permission rules (statements like "If x , y must z "), they can easily identify violations of such rules, and distinguish between intentional and accidental violations. (As in the case of language, humans are very bad at justifying their judgments, suggesting that the rules at work are part of what we have called tacit knowledge.) Along with the evidence we have concerning infants' empathy, we can safely say that quite a few of the moral capacities that characterize our species are in place very early in development, as is the case with language.

We can safely conclude, as for language, that innate biases are guiding the development of children into moral creatures (resulting in their ability to judge the moral permissibility, moral impermissibility, and moral obligatoriness of actions in situations they have never encountered before). Arguably, the capacity to distinguish between a moral rule and a conventional rule, and violations thereof, must be in place before any judgment concerning the moral permissibility or obligatoriness of an action can be made. And it is clear that children are facing a poverty of stimulus problem here too. There are countless regularities that the child could attend to, and only a few of them are characterized as falling under the moral domain. Many of the instructions provided by caretakers have exactly the same form ("you shouldn't do that!"), regardless of whether they pertain to the moral or merely conventional domain, so something internal, something biological, must guide the child in her development. That something, we may call Universal Moral Grammar,⁶ or the Moral Organ⁷.

As was also the case with language, while moral capacities are present early in development, there is some variation as to which actions are morally permissible or obligatory. This “diversity within unity” situation suggests that our moral sense is subject to parametric variation, and this possibility is entertained in recent works⁸.

Once the initial evidence for a moral organ is in place, it becomes imperative to look for the computational and structural/representational primitives that constitute moral knowledge. It also becomes important to identify the way this moral competence grows in the individual, and how it is put to use. Many details remain to be filled in, but already at this stage researchers have asked how this moral knowledge may be physically realized in the brain, and how it may have evolved in the species. In sum, they have asked the five questions at the heart of biolinguistics that I listed in Chapter 1 and that helped shape this book. Note that it may very well turn out that the specific answers to these five questions are quite different from what we find in the domain of language, but for now the most important lesson is the fruitfulness of the approach – it raises questions rarely if ever asked before, and suggests experiments to test reasonable hypotheses.⁹

The same conclusion can be reached in the domain of music. As linguist Ray Jackendoff and music theorist Fred Lerdahl made clear in their 1983 book,¹⁰ adopting a Chomskian perspective on our music capacity proves highly revealing. Jackendoff and Lerdahl stressed that what they had in mind by music capacity was not the achievements of professional musicians, just like linguists don’t have Shakespeare in mind when they discuss the creative aspect of language use. Jackendoff and Lerdahl pointed out that any human being reacts to music in specific ways (the very fact that we can pick out music from the noise around us suggests some cognitive bias). Accordingly, the first question to ask is this: When a listener hears a piece of music and recognizes it as such, what cognitive structures does she construct in response to it? What are the cognitive mechanisms involved? As Jackendoff and Lerdahl discuss in detail, when presented with a piece of music the listener spontaneously and unconsciously organizes the signal in specific ways. The listener imposes a certain grouping structure, and metrical structure, giving rise to rhythm; the listener also imposes a certain pitch structure with specific tonalities and a certain hierarchical structure (melody, theme, motif), and inevitably listening to music triggers certain emotions. All of these structures correspond to mental constructs whose rules and interactions must be uncovered by the music theorist. As Jackendoff and Lerdahl observe, within a certain musical idiom or genre, listeners are capable of understanding novel pieces of music, and judge them with remarkable accuracy and uniformity across the population, exactly as in the linguistic and moral domains.

Cross-culturally music takes different forms and idioms, so here too we are facing a diversity within unity situation that immediately begs the acquisition question: How does a listener acquire the musical grammar that she ultimately recognizes as her own, and what sort of input is required for acquisition to be successful? This question can be rephrased in more biological terms as follows: What are the pre-existing resources in the human mind/brain that make it possible for the music organ

to grow in the individual? And just as in the cases of language and morality, inquiry into the formal properties of the mental structures that inhere to the mind should run in parallel with experimental research regarding the processing of music in real time. And here too, we would ultimately like to know how the music organ is implemented in neural terms, and how it emerged in the species. As for the latter question, Jackendoff and Lerdahl have correctly pointed out that the specificity question is particularly interesting: What aspects of our musical competence are consequences of more general cognitive capacities, and what aspects are unique to music? As many have pointed out since Alfred Russell Wallace (the co-founder of evolutionary theory with Darwin), language, mathematics, and music appear to be distinctly human traits (though no doubt they built on older cognitive resources), and many have tried to relate these three aspects and find a common core to them. Building on the Hauser, Chomsky and Fitch¹¹ study at the core of Chapter 11 of this book, Jackendoff and Lerdahl¹² have argued for the usefulness of the distinction between a music faculty in the broad sense and a music faculty in the narrow sense in an attempt to determine the cognitive overlap between music and especially language (more precisely, phonology/sound structure), and research is underway to figure out what to include in the narrow set or broad set.

As can be seen, music appears to be another cognitive domain where substantial progress can be made if we ask: For a given cognitive capacity X:

- What are the rules and structures involved in X that are necessary to account for what we tacitly know about X?
- How does that capacity develop in the individual?
- How is X put to use?
- How is X implemented in the brain?
- How did X evolve?

These ethology-inspired questions – which form the basis of what one might call bio-cognition – have guided us throughout this book, and judging from the results reported on here one can say that without a doubt they provide a particularly useful “grid” with which to approach mental phenomena. This sort of methodology certainly proved invaluable in the language domain, and there are some signs it will bear the same fruits in the context of music and morality.

As Noam Chomsky pointed out:¹³

There is reason to believe that knowledge of language, which provides an unbounded range of propositional knowledge and enters into complex practical knowledge, should be regarded as a system of principles that develops in the mind by fixing values for certain parameters on the basis of experience, yielding systems that appear to be highly diverse but that are fundamentally alike in deeper respects . . . We might speculate that the same is true in other areas where humans are capable of acquiring rich and highly articulated systems of knowledge under the triggering and shaping effect of experience.

So, in the end, my hope is that the main message the reader will take away from this book is that modern cognitive science has worked out over the past 50 years a terrific point of entry into the mind and its relation to the brain. Whatever results the reader of this book encounters, she has all the tools to be able to interpret these results, and place them in a very rich intellectual context. The foundations are quite solid, and remember: the best is still to come.

Notes

Prologue

- 1 On psychology as “the science of mental life,” see William James 1890: 1.
- 2 Davis 1997 (and many subsequent reprises by cognitive scientists and neuroscientists). Consider the fact that the brain contains billions of neurons connected by trillions of synapses.
- 3 See O’Neil 1969; Keyser 1970; C. Chomsky et al. 1985; Fabb 1985; Carey 1986; Carey et al. 1986; Honda and O’Neil 1993, 2008.
- 4 Larson 1996.
- 5 Jackendoff 2007. The points raised in this paragraph summarize the discussion in a talk given by Ray Jackendoff at a conference on learning and the brain, Cambridge, November 5–8, 2003. The title of the talk was “The Structure of Language: Why it matters to education.”

Chapter 1 Mind Matters: Chomsky’s Dangerous Idea

- 1 For more information, go to Harvard’s virtual tour: www.hno.harvard.edu/tour.
- 2 Voltaire 1756, chap. 70.
- 3 For an accessible discussion, see Chomsky 1975, chap. 1, “On Cognitive Capacity.” The relevant passage is on p. 10:

Human cognitive systems, when seriously investigated, prove to be no less marvelous and intricate than the physical structures that develop in the life of the organism. Why, then, should we not study the acquisition of a cognitive structure such as language more or less as we study some complex bodily organ?

Consider also this more recent passage (from Chomsky 2004: 380):

The faculty of language can reasonably be regarded as a “language organ” in the sense in which scientists speak of the visual system, or immune system, or circulatory system, as

organs of the body. Understood in this way, an organ is not something that can be removed from the body, leaving the rest intact. It is a subsystem of a more complex structure. We hope to understand the full complexity by investigating parts that have distinctive characteristics, and their interactions. Study of the faculty of language proceeds in the same way.

- On mental organs and “cognitive physiology,” see Anderson and Lightfoot 2002, especially chap. 10.
- 4 On biolinguistics, see Jenkins 2000; see also Boeckx and Grohmann 2007, and more generally www.biolinguistics.eu.
 - 5 Dennett 1995.
 - 6 Richard Dawkins, who has written eloquently on why Darwinism still requires lengthy defenses (see Dawkins’ new preface to the 1996 edition of *The Blind Watchmaker*) once remarked that being forced to defend the most basic tenets of Darwinism whenever one is writing on evolution (and I’d add, our language faculty) is a bit as if writers on baseball had to introduce the rules of baseball every time they write a piece about America’s favorite pastime. Fortunately, baseball writers don’t have to do so, but unfortunately linguists and evolutionary biologists do.
 - 7 In the case of Darwin, the consequences for human nature implicit in *On the Origin of Species* (Darwin 1859) were made explicit in *The Descent of Man* (Darwin 1871); in the case of Chomsky, what was implicit in *Syntactic Structures* (Chomsky 1957) became explicit in Chomsky’s review of Skinner’s book *Verbal Behavior* (Chomsky 1959) and in *Aspects of the Theory of Syntax* (Chomsky 1965).
 - 8 Leibniz 1982: 3.7.6.
 - 9 James 1890: 1.
 - 10 The term “Science of Man” was very popular in philosophical circles around David Hume (the period often called the Scottish Enlightenment). Hume himself used the phrase in his *Treatise of Human Nature* (1739). For a good overview of the original Science of Man project, see Jones 1989.
 - 11 Descartes 1985.
 - 12 Chomsky 1986.
 - 13 This sentence was first examined in Chomsky 1965.
 - 14 Other favorite examples in linguistics texts are: “One morning I shot an elephant in my pajamas” (Groucho Marx), or “Dr. Ruth will discuss sex with Dick Cavett.”
 - 15 Halle 1978.
 - 16 For many more examples of this kind, see Hoffman 1998.
 - 17 I owe this parallelism between the linguistic and the visual domains to Lightfoot 1982: 45.
 - 18 See again Hoffman 1998. For a discussion of this sort of mental construct in the context of language and cognitive science, see Jackendoff 1994: 173.
 - 19 For a very good discussion of prescriptive vs. descriptive grammar, see Pullum 2004.
 - 20 On FLN and FLB, see Hauser, Chomsky, and Fitch 2002. A similar distinction is suggested (though conceived of somewhat differently) in Jackendoff 1987.
 - 21 The examples in (6–11) are part of a domain of inquiry linguists call Binding. For increasingly sophisticated overviews of binding, see Carnie 2007, chap. 5; Haegeman 1994, chap. 4; and Büring 2005.
 - 22 Chomsky 1965, chap. 1; 1966.
 - 23 Jerry Fodor, one of the great architects of modern cognitive science, has made this point on numerous occasions; see especially Fodor 2000 and (2008, chap. 4).

- 24 “There’s a reason physicists are so successful with what they do, and that is they study the hydrogen atom and the helium ion and then they stop” (quoted in Rigden 2002: 43). For a recent discussion on how novelists may anticipate results that later become the objects of scientific inquiry, see Lehrer 2007.
- 25 Descartes 1985.
- 26 Chomsky on problems and mysteries, 1975, Part II; see also Chomsky 1991, 2008, and McGinn 1989.
- 27 The five questions are articulated in accessible form in Chomsky 1986 (which focuses on the first three questions) and Chomsky 1988.
- 28 Lerdahl and Jackendoff 1983; Jackendoff and Lerdahl 2006.
- 29 Mikhail 2000, 2007; Dwyer 1999, 2006; Hauser 2006, who all follow a suggestion in Rawls 1971.
- 30 The exact quote is: “The object of the idea constituting the human Mind is the Body” (Ethics, Proposition 13, Part 2).
- 31 Pinker 1997.
- 32 See Poeppel 2005; Embick and Poeppel 2006.
- 33 See, e.g., Miller 2003. Chomsky himself would prefer a phrase like “the cognitive revival,” since he views modern cognitive science as firmly rooted in the original cognitive revolution that took place in the seventeenth century.
- 34 Gardner 1985.

Chapter 2 The Mechanization of the Mind Picture

- 1 “Those who cannot remember the past are condemned to repeat it.” George Santayana; *The Life of Reason (1905–1906)*, vol. I: *Reason in Common Sense*.
- 2 My favorite overview is, unfortunately, not (yet) available in English; it is written in Italian, by Piattelli-Palmarini 2008. My presentation of the cognitive revolution draws heavily on Chomsky’s writings on the topic, especially, Chomsky 1968, chap. 1 “Linguistic contributions to the study of mind: Past.” Miller 2003 is a brief overview that nicely complements Chomsky’s own portrayal. Boden’s 2006 massive two-volume set is very thorough, but unfortunately grossly misrepresents the role of linguistics, in a way that may not be apparent to the novice but that becomes clear very quickly to anyone somewhat familiar with the material (as Chomsky 2007b has pointed out). Nevertheless, it contains relevant material, especially in relation to the development of Artificial Intelligence. See also Gardner 1985.
- 3 See Chomsky 1955, 1957, 1959, 1965, 1966.
- 4 The term “computational” will become clear as the reader proceeds in this chapter; the term “representational” will not be clarified to the same extent, as I am not convinced it is equally adequate. Philosophers and cognitive scientists often use the term “representation.” If used in the sense of “structure” or “symbol” then the term is appropriate. But if the term is meant to indicate some mental object representing aspects of the world/environment, it runs into the sort of difficulties discussed in Chapter 7.
- 5 Subtitle of Hume’s *Treatise on human nature* (1739–40).
- 6 Skinner 1957.
- 7 Here I would like to highlight the work of the Gestalt psychologists (Max Wertheimer, Kurt Koffka, Wolfgang Kohler) working in Germany in the early decades of the last century, and focusing on visual perception.

- 8 See Skinner 1938, 1953; Watson 1919, 1924.
- 9 Pavlov 1927.
- 10 The example of Newton's work is a favorite of Chomsky's (see Chomsky 2000, especially chap. 4; and Chomsky 2002, chap. 3). Chomsky cites the example of Newton in all his discussions of the mind–body problem and the specter of Descartes' dualism. Of immediate concern is the fact that Newton went on to posit some entity whose physical nature was not known at the time and in fact conflicted with the then dominant notion of what people were willing to posit. Numerous examples from which a similar lesson could be drawn can be found in the history of science.
- 11 Lashley 1951.
- 12 Published in part in 1975.
- 13 Chomsky 1957.
- 14 For the clearest statement of this position, see Chomsky 1986.
- 15 On the remarkably linear character of the development of modern linguistics, see Freidin and Vergnaud 2001; Boeckx 2006; and Boeckx and Hornstein 2007.
- 16 Pinker 2002.
- 17 This is particularly true in cognitive neuroscience, as Gallistel and King (2009) demonstrate in detail.
- 18 Massimo Piattelli-Palmarini, personal communication.
- 19 Bloom 2004.
- 20 Piattelli-Palmarini 1994.
- 21 On more illusions of this type, see Hoffman 1998.
- 22 Lenneberg 1967; see also Lenneberg 1960, 1964, 1966.
- 23 Tinbergen 1951. See also Eibl-Eibesfeldt 1970.
- 24 Chomsky 1975.
- 25 Chomsky 1965 and subsequent works.
- 26 James 1887.
- 27 Eibl-Eibesfeldt 1970, 65f. and references therein.
- 28 Chomsky 1959.
- 29 Most of Lorenz's technical works appeared in German well before they were made available in English; for a good collection, see Lorenz 1970, 1971.
- 30 For more on "instinct–learning intercalation" and "deprivation experiments," see Eibl-Eibesfeldt 1970.
- 31 For a variety of examples, see Eibl-Eibesfeldt 1970.
- 32 Anderson 2004 on "subsong," and references therein.
- 33 Uexküll 1921.
- 34 Lorenz 1939.
- 35 Examples taken from Eibl-Eibesfeldt 1970.
- 36 Eibl-Eibesfeldt 1970.
- 37 Lenneberg 1960, 1964, 1967.
- 38 See the last chapter of Eibl-Eibesfeldt 1970 entitled "The ethology of man," and the massive "Human ethology" sequel, Eibl-Eibesfeldt 1989.
- 39 Darwin 1871.
- 40 Tinbergen 1963; see also Lorenz 1937.
- 41 Chomsky 1986, 1988.
- 42 Pinker 1994.
- 43 Anderson and Lightfoot 2002.

- 44 Hauser 2006, chap. 4.
- 45 Gallistel 2007; see also Gallistel and King 2009.
- 46 See, e.g., Chomsky 2000.
- 47 On the math organ, what Dehaene calls the number sense, see Dehaene 1997; Butterworth 1999.
- 48 Gallistel 1990. See also the second part of Gallistel and King 2009.
- 49 Chap. 1 of Eibl-Eibesfeldt 1970.
- 50 Pernau 1716.
- 51 Reimarus 1762.
- 52 Witness: Damasio's *Descartes' Error* (1994), and *Looking for Spinoza* (2003); Fodor's *Hume Variations* (2003).
- 53 Dijksterhuis 1961.
- 54 For explorations of the prehistory of modern cognitive science, see Brook 2007. See also Macnamara 1999.
- 55 All the passages I quote in this section of the book can be found in Chomsky 1966. Nevertheless, I decided to indicate the original source as well. The reader can find this first quote on p. 51 of Chomsky 1966 (second edition), and it is from Descartes (1985, vol. I: 39).
- 56 Chomsky 1966: 51; Descartes 1985, vol. I: 39–40.
- 57 Chomsky 1966: 52; Descartes 1985, vol. I: 40.
- 58 Chomsky 1966: 52.
- 59 On Vaucanson, see Boden 2006, chap. 2. On Descartes and machines, see Wheeler 2008.
- 60 For a very handy volume combining Descartes' original discussion, as well as Turing's own proposal, see Shieber 2004. I strongly recommend Chomsky's contribution in Shieber's book.
- 61 See especially Pettit 2008 and Dascal 2007.
- 62 Brook 2007.
- 63 See various contributions in Brook 2007; on Hobbes, see also Pettit 2008; on Hume, see Fodor 2003.
- 64 For an accessible discussion of this basic opposition, see Hornstein 2005.
- 65 Hume 1748.
- 66 Fodor 2003.
- 67 See chap. 2 of Chomsky 1966. This part can be somewhat challenging, as it relies on concepts like deep structure and surface structure which won't be discussed in this book, except indirectly in Chapter 4.
- 68 On this, see Chomsky 1965, chap. 1.
- 69 Chomsky 1966: 94; Cherbury 1624: 126, 133.
- 70 Chomsky 1966: 95; Cherbury 1624: 105–6.
- 71 Descartes 1641.
- 72 Chomsky 1966: 99–100; Descartes 1985, vol. I: 304–5.
- 73 Cudworth 1688 [1996].
- 74 Specifically, Plato's theory of metempsychosis (reminiscence).
- 75 Leibniz 1686, "Discourse on metaphysics."
- 76 Darwin, Notebooks.
- 77 Chomsky 1966: 103.
- 78 See the introduction in Brook 2007; see also Dascal 2007, and Pettit 2008.

- 79 See also Fodor 1975, 2008, who rightly takes “computationalism” to be the best thing that ever happened to cognitive science.
- 80 Marcus 2001.
- 81 See Fodor and Pylyshyn 1988; Pinker 1997, 2002; Marcus 2001; and Gallistel and King 2009.
- 82 The contrast is discussed in Chomsky 1957.
- 83 The contrast is discussed in Chomsky 1975.
- 84 Pettit 2008 sees this as the core of Hobbes’ philosophy. I encourage readers who want to find out more about this central aspect of Hobbes’ thought to turn to Pettit’s book, as well as to Dascal 2007, for quotes and relevant discussion.
- 85 For relevant discussion, see Viertel 1966 and Schreyer 1985.
- 86 Chomsky 2002, chap. 3 contains many relevant quotes, and much relevant discussion; see also Yolton 1983.
- 87 Wheeler 2008.
- 88 Hobbes 1651.
- 89 See Boden 2006, and some of the contributions in Husbands, Holland, and Wheeler 2008.
- 90 For an accessible historical overview, see Davis 2000; most of the seminal papers in this area are collected in Davis 2004.
- 91 See Turing 1936. For an excellent discussion of computation, which I have relied on to write this part of the book, see Gallistel and King 2009.
- 92 On Hilbert’s challenge, see Gray 2000. Hilbert’s program refers to works tackling a list of twenty-three problems in mathematics put forward by German mathematician David Hilbert at the Paris conference of the International Congress of Mathematicians in 1900.
- 93 Turing 1936.
- 94 I owe this example and the following discussion to Gallistel and King, 2009.
- 95 “Language is really situated in relation to an infinite and truly boundless sphere, the epitome of everything that can be thought. Thus it must make an infinite use of finite means and is able to do so through the identity of the faculty that generates thoughts and language” (Humboldt 1827).
- 96 For a very clear statement, and related discussion, see The Church-Turing thesis entry in the Stanford Encyclopedia of Philosophy, <http://plato.stanford.edu/entries/church-turing/>.
- 97 Chomsky 1956, 1957.
- 98 This is a favorite theme of Chomsky’s works, as we already saw. For a very good discussion of Descartes’ relation to machines, see Wheeler 2008.
- 99 On this point, see Hodges 2008.
- 100 Fodor 2000, reviewing Pinker 1997.
- 101 See Dijksterhuis 1961; see also Chomsky 2000, 2002, in press.
- 102 See Shannon 1948; Shannon and Weaver 1949.
- 103 Chomsky 1957; see also Miller 2003.
- 104 Gallistel 2001, 2006, 2007, 2008; it is also a major theme of Gallistel and King 2009.
- 105 I owe the example of the thermometer to Gallistel 2001.
- 106 I owe this example and the following discussion to Nevins 2004.
- 107 Chomsky 1968: 80–1; see also Chomsky 1980, 2000; see Peirce 1931–66: 237ff.
- 108 Nevins 2004.

- 109 On selective vs. constructive grammars, see Piattelli-Palmarini 1987, 1989; see also Moro 2008 (who talks about the difference between selective and constructive recipe-modes), and, in the context of immunology, Jerne 1985.
- 110 Nevins 2004.
- 111 Marr 1982.

Chapter 3 How the Mind Grows: From Meno to Noam

- 1 Quine 1960.
- 2 Locke 1690. For valuable discussion, see also Gleitman 1990; Bloom 2000; and Gleitman and Fisher 2005.
- 3 On “biased embryos,” see Arthur 2004.
- 4 Plato [1956] or equivalent editions.
- 5 On the wug test, Gleason 1958.
- 6 Chomsky 1980; see already Chomsky 1965: 55–6, and Chomsky 1968: 61–3 on “structure-dependence,” and Chomsky 1971, and various portions of Chomsky’s contribution in Piattelli-Palmarini 1980.
- 7 For an excellent collection on pidgin and creoles, see DeGraff 1999. Bickerton’s work (such as Bickerton 1981) forms the starting point of most current views on how creoles emerge.
- 8 Bickerton 1981, 1984; the so-called bioprogram hypothesis.
- 9 Descartes 1641.
- 10 Goldin-Meadow 2003.
- 11 Among many relevant discussions, see Kegl, Senghas, and Coppola 1999; Kegl 2004; Saffran, Senghas, and Trueswell 2001.
- 12 Sandler, Meir, Padden, and Aronoff 2005; Senghas 2005; and for a non-technical overview, Fox 2007.
- 13 For excellent surveys of cases of language creation, see Petitto 2005 and Kegl 2004.
- 14 Ochs 1982.
- 15 McNeil 1966.
- 16 Chomsky 1965, 1968.
- 17 Gleitman and Landau 1985.
- 18 Lenneberg 1960, 1964, 1966.
- 19 Lenneberg 1967.
- 20 See Eibl-Eibesfeldt 1970 and references therein.
- 21 Hubel and Wiesel 1962 and subsequent works; collected in Hubel and Wiesel 2004.
- 22 Mehler 1974.
- 23 Changeux, Courrège, and Danchin 1973; Edelman 1978, 1989.
- 24 For accessible treatments of brain growth, see Dowling 1998, 2004; Marcus 2004.
- 25 Gallistel and King, 2009.
- 26 Some of the key experiments go back to the 1970s; see Lasky, Syrdal-Lasky, and Klein 1975; Trehub 1976; Streeter 1976; see also Aslin, Pisoni, Hennessy, and Perey 1981. For a comprehensive treatment of phonological development, see Jusczyk 1997.
- 27 Eimas, Siqueland, Jusczyk and Vigorito 1971.
- 28 DeCasper and Fifer 1980.
- 29 DeCasper and Spence 1986.

- 30 Friedlander 1970.
- 31 Eisenberg 1975.
- 32 For good overviews, see Mehler and Dupoux 1994 and Jusczyk 1997.
- 33 Moon, Cooper, and Fifer 1993.
- 34 For additional feats, see Vouloumanos and Werker 2007 and references therein.
- 35 Vouloumanos and Werker 2004.
- 36 For a textbook introduction to the various stages, see Crain and Lillo-Martin 1999.
- 37 Petitto et al. 2001.
- 38 See Sandler and Lillo-Martin 2006.
- 39 On feral children, see Malson 1972; Lane 1976; Newton 2002; Benzaquen 2006.
- 40 See Curtiss et al. 1974; Fromkin et al. 1974; Curtiss 1977; and Rymer 1994.
- 41 Kingsley 1940, 1947.
- 42 See Keyser 2001.
- 43 See Lenneberg 1966.
- 44 Lenneberg 1967; see also Kegl's felicitous phrase "language-ready brain" (Kegl 2004).
- 45 McCloskey 2004.
- 46 Lenneberg 1967.
- 47 West-Eberhard 2003.
- 48 See Lewontin 2000a.
- 49 See also Lewontin 2000a, 2000b.

Chapter 4 Mental Chemistry

- 1 Wilczek and Devine 1987.
- 2 Hence the allusion to Keats' famous poem *Lamia* about Newton unweaving the rainbow (Keats 1819), used as the title of a book by Dawkins (1998) to celebrate the poetry of scientific discovery.
- 3 Moro 2008.
- 4 Carroll 1872.
- 5 Named after Gaetano Kanizsa, who discussed such optical illusions in 1955; for more, see Hoffman 1998.
- 6 Gleitman 1990; Gleitman and Fisher 2005, and references therein.
- 7 On levels of representation, see Chomsky 1955, and for an accessible overview, Lasnik 2005.
- 8 The discussion that follows was inspired by, and draws heavily from, Lasnik 2000, which is a commentary on Chomsky 1957. Portions of my presentation are adapted from Boeckx 2006.
- 9 Chomsky 1986.
- 10 Chomsky 1957.
- 11 The material also lies in the immediate background of Gallistel and King, 2009, who use many of Chomsky's arguments to show how cognitive science may transform neuroscience.
- 12 For an excellent discussion of the desirability (and necessity) of formalism, see Moro 2008.
- 13 Chomsky 1957; see also Chomsky 1956.
- 14 Lasnik 2000.

- 15 I owe this example to Uriagereka 2008.
- 16 Chomsky 1957.
- 17 If the reader is interested in learning more, she should turn to Lasnik 2000.
- 18 The term “mental chemistry” was introduced by John Stuart Mill in 1843.
- 19 On fractals, see Mandelbrot 1982; see also the very useful website at: <http://classes.yale.edu/Fractals>.
- 20 See the references at <http://classes.yale.edu/Fractals>.
- 21 On endocentricity, see Chomsky 1970. For an accessible discussion, see Hornstein, Nunes, and Grohmann 2006, who introduced the term “periscope.”
- 22 For introductory discussions, see Carnie 2007 and especially Haegeman 2006. See also Pollock 1989; Cinque 1999; Rizzi 1997; and Boeckx 2008 for more challenging material.
- 23 Weinberg 1993.
- 24 The result was published in Chomsky 1970; see also Jackendoff 1977.
- 25 The discussion to follow draws heavily on Boeckx 2006.
- 26 See Boeckx 2009, Hornstein 2009; on “headed” structures, see also Jackendoff 1987.
- 27 See Carnie 2007; see also Lasnik 2000.
- 28 Chomsky 1965, 1968; see also the poverty of stimulus argument in Chapter 3.
- 29 See Kayne 1994 and much subsequent work, among which Moro 2000 (introduced in accessible form in Moro 2008, chap. 3).
- 30 For an accessible overview, see Randall 2005.
- 31 First developed by Desargues in the seventeenth century, it did not achieve prominence as a field of mathematics until the early nineteenth century. Projective geometry originated from the principles of perspective art.
- 32 Marr 1982.
- 33 Pizlo 2008.
- 34 Again, see Randall 2005 for an accessible discussion.
- 35 All of whom take Kayne 1994 as their starting point. For an excellent overview of Kayne’s system, see Hornstein, Nunes, and Grohmann 2006.
- 36 Locality is one of the most important domains of research in linguistics. Locality issues were first discussed in Chomsky 1964 and Ross 1967. For a very limited sample of seminal works in this area, see Chomsky 1973; Chomsky 1977; Koster 1978; Kayne 1984; Rizzi 1990; and Cinque 1990. All of these works are challenging. For introductory treatments, see Haegeman 1991 and Adger 2003.
- 37 Example and discussion adapted from Crain and Lillo-Martin 1999.
- 38 See Rizzi 1990; Chomsky and Lasnik 1993; Chomsky 1995; Boeckx 2006; Lasnik, Uriagereka, and Boeckx 2005. For intriguing discussion of why this might be, see Fukui 1996, Uriagereka 1998, and Chomsky 2005.
- 39 For modern appreciations of the works of Gestalt psychologists, see Neisser 1967; Marr 1982; Pizlo 2008; and virtually every work by Ray Jackendoff (especially Jackendoff 1983, and Lerdahl and Jackendoff 1983).
- 40 This constraint is known in linguistics as the Coordinate Structure Constraint; see Ross 1967.
- 41 On traces left by displacement operations, see Chomsky 1976, 1981; Fiengo 1977.
- 42 On copies left by displacement operations, see Chomsky 1993.
- 43 For a popular introduction, see Pickover 2006.
- 44 This example is adapted from Crain and Lillo-Martin 1999.

- 45 An adaptation of a felicitous statement due to French physicist Jean Perrin (cited in Dehaene 2007). The original is “explain the complexities of the visible in terms of invisible simplicity.”
- 46 For a discussion of simplicity at all these levels, see Boeckx 2006 and Moro 2008.
- 47 Chomsky 1965. For more on the competence–performance distinction, see Chapter 9.
- 48 The same point is made in Lasnik 2000.
- 49 Cited in Smith 2005.
- 50 Marr 1982.
- 51 This approach is known as the minimalist program; for an accessible introduction, see Boeckx 2006.
- 52 Miller 1956.
- 53 Chomsky is fond of using this analogy to convey the message in the text; see Chomsky 2000, 2002, in press.

Chapter 5 The Variety of Linguistic Experience: The Towers of Babel and Pisa

- 1 Gordon 2005.
- 2 Cf. Max Weinreich’s aphorism, “A language is a dialect with an army and navy” (Weinreich 1945).
- 3 The great linguist Richard Kayne has made this point on numerous occasions; see Kayne 2000, 2005.
- 4 Kayne 2000.
- 5 Baker 2001.
- 6 The watershed is Chomsky’s 1979 Pisa lectures (published as Chomsky 1981).
- 7 The 1979 annual meeting of GLOW (Generative Linguistics in the Old World; Koster, van Riemsdijk, and Vergnaud 1978).
- 8 Chomsky was of course not working in a vacuum; crucial components that led to his proposal were Chomsky 1973; Kayne 1975; and Rizzi 1978.
- 9 The introductory chapter of Chomsky 1981 is accessible, and well worth reading; see also Chomsky 1986, 1988; and Baker 2001. Roberts 2006 contains all the major papers on parameters with very useful introductory comments.
- 10 See Lightfoot 1982, 1991, 1999, 2006.
- 11 The metaphor is attributed to Jim Higginbotham in Chomsky 1986.
- 12 For many more examples, see Baker 2001.
- 13 Huang 1982.
- 14 Huang 1982; Lasnik and Saito 1984; Watanabe 2001.
- 15 Rizzi 1986 and much subsequent work.
- 16 Emonds 1978; Pollock 1989; Cinque 1999.
- 17 Chomsky 1986.
- 18 Greenberg 1966.
- 19 Chomsky 1981.
- 20 See Baker 2001, and for a more technical overview, Cinque and Kayne 2005.
- 21 Jacob and Monod 1961. Chomsky 1980 acknowledges the parallelism.
- 22 For the most accessible treatment that I know of, see Carroll 2005.
- 23 Arthur 2004.

- 24 See Carroll 2005 for relevant discussion on the biology side; for linguistics, see Joos 1957.
- 25 This is the so-called multiple wh-fronting phenomenon; for an overview, see Boeckx and Grohmann 2003.
- 26 Baker 2001.
- 27 The specific hierarchy is reproduced from Baker 2003. For additional discussion of the hierarchy, see Baker 2005, and Newmeyer 2005.
- 28 Baker 2001.
- 29 See Campbell and Poser 2008.
- 30 See Rizzi 1989 and Roberts 2006 (introduction). For discussion of how results of the two modes of comparison may converge, see Longobardi 2003.
- 31 Halle and Richards 2007, on which the text discussion heavily draws.
- 32 See Hale and Keyser 1993, 2002.
- 33 Saussure 1916.
- 34 For an accessible description, see Baker 2001; for a much more technical treatment, see Baker 1996.
- 35 Lerdahl and Jackendoff 1983 and Jackendoff 1989 for music; Hauser 2006; Dwyer 1999, 2006; Mikhail 2000, to appear.

Chapter 6 All Roads Lead to Universal Grammar

- 1 See DeGraff 1999 for a volume devoted to language creation and language change; the relation between the two phenomena is a constant theme of David Lightfoot's work (see Lightfoot 1982, 1991, 1999, 2006) and Tony Kroch's (see Kroch 2000). See also Crain, Goro, and Thornton 2006; and Niyogi 2006.
- 2 Howard Lasnik, personal communication.
- 3 Meroni, Gualmini, and Crain 2001; Crain and Thornton 1991, 1996.
- 4 Thornton 1995, 2008; Thornton and Tesan 2007; Tesan and Thornton 2003, 2005.
- 5 I owe this analogy to Meroni, Gualmini, and Crain 2001.
- 6 This point is emphasized in Yang 2002, 2005; and Pearl 2007.
- 7 For a representative sample, see Berwick 1985, Wexler and Manzini 1987; Clark and Roberts 1993; Gibson and Wexler 1994; Berwick and Niyogi 1996; Fodor 1998, 2001; Sakas and Fodor 2001; Yang 2002; Pearl 2007; Lidz and Pearl 2007.
- 8 Bloom 2000.
- 9 A distinction emphasized in Pearl 2007.
- 10 See Gallistel and King, 2009.
- 11 A point stressed by Yang 2002.
- 12 Yang 2005.
- 13 See especially Pearl 2007.
- 14 Drescher and Kaye 1990; Drescher 1999; Lightfoot 1999; Pearl 2007.
- 15 Yang 2002.
- 16 Lightfoot 1989.
- 17 Borer and Wexler 1987, 1992; Wexler 1999.
- 18 Often called Bayesian, after Thomas Bayes (1702–61). On the recent resurgence of interest in probabilistic learning mechanisms, see Chater, Tenenbaum, and Yuille 2006.
- 19 See Yang 2004 (on which this section of the text heavily relies) on this point.

- 20 Yang 2002, 2004, 2006.
- 21 Mayr 1970, 1982.
- 22 Witness Elman, Bates, Johnson, Karmiloff-Smith, Parisi, and Plunkett 1997.
- 23 Drawing on works by Gallistel; see Gallistel and King, 2009.
- 24 See Yang 2004, on which I draw heavily in the text that follows; see also Gambell and Yang 2004a and 2004b.
- 25 Saffran, Aslin, and Newport 1996.
- 26 Gambell and Yang 2004a and 2004b.
- 27 Morgan and Demuth 1996; Christophe et al. 1997; Ramus 1999; Ramus and Mehler 1999; Nespor 2001.
- 28 Sapir 1921.
- 29 Yang 2005. For a different treatment of the productivity question, see Pinker 1999.
- 30 Developing ideas of Morris Halle (e.g., Halle 1990).
- 31 I extend this logic to other aspects of language learning in Boeckx to appear a.
- 32 Marcus et al. 1992.
- 33 On cues, Dresher and Kaye 1990; Dresher 1999; Lightfoot 1999. The term “signature” is due to Yang 2002.
- 34 See Lightfoot 1999.
- 35 Hubel and Wiesel 1962; Sperry 1968; Hubel 1988.
- 36 Lightfoot 1989, departing from less restrictive proposals in Wexler and Culicover 1980.
- 37 Fodor 1998.
- 38 For an excellent comparison of cue-based approaches and Fodor’s model, see Pearl 2007, which influenced this part of the text.
- 39 This is the point made in Legate and Yang 2002, addressing argument by Pullum and Scholz 2002.
- 40 This portion of the text relies on material discussed in Boeckx and Hornstein 2007, which in turns draws heavily on Legate and Yang 2002.
- 41 This section of the text was influence by the presentation of the lessons in Meroni, Gualmini, and Crain 2001. For many aspects of child language, see Guasti 2002.
- 42 Pinker 1984; Crain 1991; Goodluck 1991; Crain and Pietroski 2001, 2002.
- 43 See Thornton 1990, 2008.
- 44 Thornton 2008.
- 45 Rizzi 2001.
- 46 Thornton 1990, 1995.
- 47 McDaniel 1986.
- 48 See Piattelli-Palmarini 1989.
- 49 Piaget 1928, 1952, 1955; for an overview (and a clear demonstration of the limitations of this approach), see Piattelli-Palmarini 1980.
- 50 See the sections by Chomsky in Piattelli-Palmarini 1980.
- 51 Dresher 2003; see also Baker 2005.

Chapter 7 Making Sense of Meaning: An Instruction Manual

- 1 Jackendoff 2002.
- 2 An example often used by Chomsky; see Chomsky 1968, 2000.
- 3 Dresher 2003 makes the interesting point that although humans automatically know how to narrow down the search for possible meanings given a certain signal (cf. the

Gavagai story in Chapter 3), it is much harder for humans (even experts) to determine what the meanings of the signals of, say, vervet monkey calls may be. Drescher is right to see in this a reflex of our biology.

- 4 Cf. Barbara Partee's question in the title of her 1979 paper "Semantics – Mathematics or Psychology?"
- 5 See, e.g., Heim and Kratzer 1998, chap. 1.
- 6 Davidson 1967a.
- 7 Montague 1970a, 1970b, 1973.
- 8 For various textbook treatments (listed in order of increasing complexity), see Portner 2005; Chierchia and McConnell-Ginet 1990; Larson and Segal 1995; and Heim and Kratzer 1998.
- 9 A statement adapted from Pietroski 2003, one of the many papers of Pietroski's on which I rely heavily in this chapter.
- 10 Jackendoff 1992, building on Chomsky 1986.
- 11 On the importance of "negative" facts, see especially Pietroski 2005a.
- 12 For a few accessible works, see Jackendoff 1983, 1990; Hornstein 1984; Moravcsik 1998; McGilvray 1999; Bilgrami 2002; Pietroski 2003, 2005a, 2006; Stainton 2006. For more technical treatments, see Pustejovski 1995; Pietroski 2005b, 2007, to appear; Hinzen 2007; and Uriagereka 2008.
- 13 Chomsky frequently quotes Strawson's 1950 claim that words don't refer, but people refer by using words.
- 14 The meaning of natural-kind terms like *gold* or *water* and technical terms like *arthritis* are favorite themes of advocates of E-semantics (Putnam 1975, Burge 1979); for discussion, see Chomsky 2000.
- 15 Pietroski 2005a.
- 16 Chomsky 1965.
- 17 A sentence discussed in Higginbotham 1995, based on the sentence *I had a book stolen* discussed in Chomsky 1965.
- 18 Example based on Pietroski 2005a, who discusses *The millionaire called the senator from Texas*.
- 19 Wittgenstein 1953.
- 20 Many of the examples of negative facts discussed in the texts are due to Pietroski 2005a.
- 21 Example inspired by Donald Davidson's famous example *Jones buttered the toast in the bathroom at midnight* (Davidson 1967b).
- 22 Expressions like *the vase broke* are known as "middle constructions" in the linguistics literature.
- 23 This latter implicational property of so-called quantifiers like *every* is known in the literature as "conservativity" (see Barwise and Cooper 1981; Higginbotham and May 1981; and Keenan and Stavi 1986).
- 24 My brief excursus on "negative polarity" draws on Portner 2005; the key findings reported on are due to Fauconnier 1975 and Ladusaw 1979. The roots of the role of the directionality of entailment go back to medieval logicians; see Ludlow 2002.
- 25 See Larson and Segal 1995.
- 26 Portner 2005.
- 27 Pietroski 2005a.
- 28 Chomsky 1965. Such constructions are discussed in the linguistics literature under the rubric of control.
- 29 Chomsky often characterizes linguistic semantics as syntax in the broad sense (see Chomsky 2000).

- 30 On meanings as instructions, see Chomsky 1995, 2002, who builds an analogy with the more standard take on phonology providing instructions to the articulators (tongue, lips, etc.) (Halle 1995, 2002). Paul Pietroski has explored the hypothesis in detail in recent works (see Pietroski to appear).
- 31 See Hornstein 1984.
- 32 A favorite example of Chomsky's, along with *London* (see Chomsky 2000).
- 33 I owe this example to Stainton 2006. For more examples, see Moravcsik 1998 and Pustejovsky 1995.
- 34 Examples due to Stainton 2006.
- 35 I owe the metaphor to Hornstein 2005.
- 36 Chomsky 1966 and much subsequent work.
- 37 McGilvray 1999, 2005.
- 38 Cudworth 1688 [1996].
- 39 Humboldt 1836 [1999].
- 40 Chomsky 1959.
- 41 Chomsky 2000.
- 42 McGilvray 2005.
- 43 McGilvray 2005.
- 44 On the importance of these activities for the child's cognitive development, see Carruthers 2006.
- 45 Chomsky 1965, 1966.
- 46 Chomsky 1966.
- 47 See McGilvray 2005 for excellent discussion.
- 48 See, e.g., Chomsky 2000.

Chapter 8 Wonderful Mental Life: Unthinkable in the Absence of Language

- 1 Cf. Malebranche's oft-cited passage: "They eat without pleasure, cry without pain, grow without knowing it; they desire nothing, fear nothing, know nothing" (Malebranche, *Oeuvres complètes*, 1958–78, II: 394).
- 2 See Cottingham 1978; Harrison 1992.
- 3 See Hauser 2000; Bermúdez 2003; Carruthers 2006.
- 4 See, e.g., Fellerman and van Essen 1991.
- 5 Darwin 1871.
- 6 Spelke 2003a, which provided the basis for this chapter, along with other works by Elizabeth Spelke and her collaborators.
- 7 Spelke 2003a is the best overview of this hypothesis I know of; for related proposals, see Tattersall 1998; Carruthers 2002, 2006; Chomsky 2005; Pietroski 2007; Boeckx to appear b, and (to some extent) Mithen 1996. Darwin (1871: 126) appears to point in this direction in the following passage:

If it could be proved that certain high mental powers, such as the formation of general concepts, self-consciousness, etc., were absolutely peculiar to man, which seems extremely doubtful, it is not improbable that these qualities are merely the incidental results of other highly-advanced intellectual faculties; and these again, mainly the result of the continued use of a perfect language.

- 8 One of the main lessons of Gallistel 2007, and Gallistel and King, 2009.
- 9 Fodor 1983.
- 10 Chomsky 1975.
- 11 Spelke 1988a, 1988b, 1994, 2000, 2003a, 2003b, 2004; Spelke, Breinlinger, Macomber, and Jacobson 1992; Carey and Spelke 1994, 1996; Hermer-Vasquez, Spelke, and Katsnelson 1999; Spelke and Tsivkin 2001; Santos, Hauser, and Spelke 2002; Feigenson, Dehaene, and Spelke 2004; Kinzler and Spelke 2007.
- 12 Kinzler and Spelke 2007; see also Fodor 2000.
- 13 On massive modularity, Barkow, Cosmides, and Tooby 1992; Pinker 1997, 2002; and much work in the field of evolutionary psychology; see also Carruthers 2006.
- 14 I am here following Kinzler and Spelke's 2007 assessment, as well as their basic description in the text that follows; see also Spelke 2004.
- 15 See Spelke 1990, 2003a; Wynn 1992; Baillargeon 1987, 2008, and the many references therein.
- 16 Spelke 1990, 2003a.
- 17 See, e.g., Hauser, Carey, and Hauser 2000; Feigenson, Carey, and Hauser 2002; and the experiments reported in Spelke and Kinzler 2007.
- 18 Dehaene 1997; Hauser 2000; Feigenson and Carey 2003; Hauser and Carey 2003.
- 19 Feigenson, Carey, and Hauser 2002; Feigenson and Carey 2003.
- 20 The same limits hold for monkeys; see Hauser, Carey, and Hauser 2000.
- 21 Kaufman, Lord, Reese, and Volkmann 1949.
- 22 See Dehaene 1997; Feigenson and Carey 2005; Feigenson 2007.
- 23 Dehaene 1997; Gallistel and Gelman 1992.
- 24 After Ernst Heinrich Weber (1795–1878); for an accessible discussion, see Dehaene 1997.
- 25 Xu and Spelke 2000; Lipton and Spelke 2003.
- 26 Gallistel 1990; Dehaene 1997; Hauser 2000.
- 27 Kinzler and Spelke 2007.
- 28 See Kinzler and Spelke 2007 for references.
- 29 For excellent overviews, see Gallistel 1990, and Gallistel and King, 2009.
- 30 See Kinzler and Spelke 2005, 2007; see also Bloom 2000 on theory of mind in connection with word learning, and Tomasello 1999.
- 31 Kinzler and Spelke 2007.
- 32 Spelke 2003a.
- 33 See, e.g., Spelke 2003a.
- 34 Fodor 1983.
- 35 The point that Fodor made in 1983 is emphasized in Fodor 2000.
- 36 As Fodor 1983 recognizes.
- 37 See Uttal 2001, who uses *The New Phrenology* as the title of his book.
- 38 Once again, see Gallistel and King, 2009, on this basic point.
- 39 Fodor 1983, part III.
- 40 Pylyshyn 1984, 1999.
- 41 Discussed in Fodor 1983.
- 42 Barkow, Cosmides, and Tooby 1992; Pinker 1997, 2002.
- 43 Among the most explicit in this regard: Carruthers 2002, 2006; Spelke 2003a; Pietroski 2007. For my own personal take, see Boeckx to appear b.
- 44 Much of the evidence comes from experimental work done by Spelke and her collaborators. For additional discussion, see Carruthers 2006.

- 45 These are the two domains chosen by Spelke 2003a to argue for her position. My presentation follows hers pretty closely.
- 46 See Biegler and Morris 1993, 1996.
- 47 Cheng 1986; Margules and Gallistel 1988; Gallistel 1990.
- 48 Hermer and Spelke 1994, 1996; Wang, Hermer-Vasquez, and Spelke 1999; Gouteux and Spelke 2001.
- 49 Hermer-Vasquez, Moffett, and Munkholm 2001.
- 50 Hermer-Vasquez, Spelke, and Katsnelson 1999.
- 51 Spelke 2003a discusses a fourth difference omitted here.
- 52 Hauser, Carey, and Hauser 2000; Feigenson, Carey, and Hauser 2002.
- 53 Spelke 2003a.
- 54 See, e.g., Wynn 1990.
- 55 Wynn 1990, 1992.
- 56 See O’Kane and Spelke 2001; Spelke and Tsivkin 2001.
- 57 De Villiers 2005 and references therein.
- 58 Boyer 2003.
- 59 Mithen 1996.
- 60 Hauser 2008.
- 61 Boeckx to appear b.
- 62 See also Cheney and Seyfarth 1990.

Chapter 9 Grammar Caught in the Act

- 1 Miller 1998.
- 2 Chomsky 1965.
- 3 Chomsky 1965: 3–4.
- 4 Chomsky 1965: 4.
- 5 Chomsky and Miller 1963.
- 6 Example due to Townsend and Bever 2001.
- 7 On the use and abuse of the competence–performance distinction, see Jackendoff 1992 and Marantz 2005. My position is much closer to the one advocated in Marantz’s paper, which greatly influenced the content of this chapter.
- 8 Piattelli-Palmarini 1994.
- 9 See Phillips 2004, Marantz 2005. For relevant discussion and various examples of grammatical accuracy (drawing from a large pool of experiments), see also Phillips and Lau 2004; Phillips 2006, to appear; Phillips and Wagers 2007.
- 10 Chomsky 1965: 15.
- 11 For a renewed appreciation of this point, see Marantz 2005. On integration/unification vs. reduction, see Chomsky 2000, 2002.
- 12 Marantz 2005.
- 13 See Miller and Chomsky 1963.
- 14 On this point, see also Chomsky 1965: 10ff.
- 15 Fodor, Bever, and Garrett 1974.
- 16 Miller and Chomsky 1963.
- 17 Marantz 2005.
- 18 Fodor, Bever, and Garrett 1974.
- 19 Marantz 2005.

- 20 On this point, see again Marantz 2005.
- 21 Although more care has often been called for in designing such “cheap” experiments; see Schütze 1996. For an attempt to address the growing skepticism about such cheap data in some corners of cognitive science, and an overall positive assessment of the situation, see Phillips to appear.
- 22 For some valuable suggestions, see Schütze 1996; see also Sprouse 2007.
- 23 See Ferreira 2005 and Bresnan 2007, among others.
- 24 Phillips to appear.
- 25 Alec Marantz and his students and collaborators did much to re-evaluate the alleged failure of the derivational theory of complexity. For an early reassessment, see Phillips 1996.
- 26 For the clearest statement, see Phillips 2004; see also Phillips and Lau 2004, and Phillips and Wagers 2007. For an accessible overview of case studies supporting Phillips’s conclusion, see Phillips 2006.
- 27 Taken from Phillips 2004.
- 28 For a similar conclusion, see Marantz 2005.
- 29 Stated explicitly in Hornstein 2009; also discussed in Phillips and Lau 2004.
- 30 For a very good case study, see Yoshida, Aoshima, and Phillips 2004.
- 31 For an excellent book, see Townsend and Bever 2001.
- 32 Bever 1970.
- 33 Phillips 2004.
- 34 Miller and Chomsky 1963.
- 35 Christianson, Hollingworth, Halliwell, and Ferreira, 2001; Ferreira, Christianson, and Hollingworth, 2001.
- 36 Phillips 2004.
- 37 On idiomatic expressions, see Jackendoff 1997, whose conclusions I personally resist (see Boeckx and Piattelli-Palmarini 2007).
- 38 A principle introduced in Chomsky 1981.
- 39 Example reported in Townsend and Bever 2001.
- 40 On these, see the classic Fromkin 1973.
- 41 Phillips 2004.
- 42 Marantz 2005.
- 43 This situation is slowly emerging in the study of our moral sense (see Epilogue), where judgments are being supplemented by other kinds of evidence; see Hauser 2006.

Chapter 10 The (Mis)Measure of Mind

- 1 Christiansen and Kirby 2003.
- 2 Fisher 2006.
- 3 Davis 1997.
- 4 Lewontin 1998.
- 5 Pinker 1997.
- 6 The quoted terms in this sentence were collected in Chomsky 2000, where they are discussed at length along the lines I retrace in the text. See also Chomsky 2002.
- 7 Crick 1994.
- 8 Nagel 1993.
- 9 Churchland 1994.
- 10 On this point, see, in addition to Chomsky 2000, 2002, Yolton 1983.

- 11 Priestley 1777.
- 12 See Chomsky in press for numerous quotes to this effect.
- 13 Darwin 1871.
- 14 A term used in this context by both Colin Phillips (Phillips 2004) and David Poeppel (see Poeppel 2005, Poeppel and Embick 2005, Embick and Poeppel 2006), on whose reflections I draw extensively in this part of the book.
- 15 For more detailed suggestions, see Gallistel and King, 2009.
- 16 Building on Hauser, Chomsky, and Fitch 2002.
- 17 For interesting examples and valuable discussion of the early history of neuroscience, see Finger 1994.
- 18 Broca 1861.
- 19 Wernicke 1874.
- 20 Lichtheim 1885.
- 21 Geschwind 1970.
- 22 For an excellent collection of historical documents leading up to the classic model, see Grodzinsky and Amunts 2006.
- 23 Broca 1861.
- 24 Wernicke 1874.
- 25 Lichtheim 1885.
- 26 Geschwind 1970.
- 27 Brodmann 1908, also reproduced in Grodzinsky and Amunts 2006.
- 28 Poeppel and Hickok 2004.
- 29 Poeppel and Hickok 2004.
- 30 Poeppel and Hickok 2004.
- 31 See Zurif 1980, and for a good overview, the introductory section of Grodzinsky 1990.
- 32 Poeppel and Hickok 2004 discuss all three in more detail than I can do here, and provide extensive references that reinforce their claim.
- 33 See Caramazza and Zurif 1976.
- 34 For extensive reference, see Poeppel and Hickok 2004.
- 35 See especially Hickok and Poeppel's 2000 model.
- 36 Well illustrated in the volume edited by Poeppel and Hickok 2004.
- 37 I here draw extensively on Poeppel 2005.
- 38 Uttal 2001.
- 39 Poeppel 2005.
- 40 Poeppel 2005.
- 41 For a very useful guide to these techniques, see Papanicolaou 1998.
- 42 See Phillips and Sakai 2005.
- 43 The discussion that follows is a very condensed version of the long argument presented in Moro 2008.
- 44 Musso et al. 2003.
- 45 This, I think, would be Jerry Fodor's reaction, judging from his critical reflections on brain imaging studies in Fodor 1999.
- 46 Moro et al. 2001.
- 47 References to specific studies can be found in Embick and Poeppel 2006.
- 48 Again, for specific references, see Embick and Poeppel 2006.
- 49 For relevant discussion, see Anwender et al. 2007. See also various contributions in Grodzinsky and Amunts 2006.

- 50 Poeppel 2005; see also Poeppel and Embick 2005. This portion of the text follows his reflections pretty closely.
- 51 For a discussion of these, and how they may relate to language concepts, see Pulvermüller 2003.
- 52 For an interesting attempt to do just that, in the domain of recursion, see Treves 2005, inspired by Hauser, Chomsky, and Fitch 2002. In the domain of aphasia, see the computationally explicit hypothesis in Grodzinsky 1990 and some of the contributions in Grodzinsky and Amunts 2006, especially Friedmann 2006.
- 53 See Chomsky 2000, 2002.
- 54 Ben Shalom and Poeppel 2008.
- 55 Jackendoff 2002 discusses the first three; Phillips and Lau 2004 add the fourth problem to the pool in their review of Jackendoff 2002.
- 56 Phillips and Lau 2004.
- 57 See, e.g., Chomsky 1995.
- 58 As discussed at length in Jackendoff 1997, 2002, 2007.
- 59 Marantz 2005.
- 60 Poeppel 2005.
- 61 Moro 2008.

Chapter 11 *Homo Combinans*

- 1 Société de Linguistique de Paris, Statuts 1968: iii. The London Philological Society followed suit in 1872.
- 2 Tinbergen 1963; Chomsky 1986, 1988.
- 3 By Darwinian I mean a very rich framework, much more pluralistic in nature than in many popular characterizations (such as Dawkins 1976). For an excellent model, see Gould 2002.
- 4 This will be obvious to anyone familiar with works in evolutionary psychology (Barkow, Cosmides, and Tooby 1992; Pinker 1997, 2002).
- 5 Darwin 1859. Darwin lamented this narrow conception of his theory. As he notes in the final chapter of the 6th edition of *The Origin of Species*:

But as my conclusions have lately been much misrepresented, and it has been stated that I attribute the modification of species exclusively to natural selection, I may be permitted to remark that in the first edition of this work, and subsequently, I placed in a most conspicuous position – namely, at the close of the Introduction – the following words: “I am convinced that natural selection has been the main but not the exclusive means of modification.”

This has been of no avail. Great is the power of steady misrepresentation.

- 6 On this point, see Lewontin 1998.
- 7 The point is well made in Hauser, Chomsky, and Fitch 2002, and even more forcefully in Fitch, Hauser, and Chomsky 2005. See also Piattelli-Palmarini 1989; Uriagereka 1998; and Lorenzo and Longa 2003.
- 8 Chomsky has made this point very forcefully in recent years (see Chomsky 2002, 2005; Hauser, Chomsky, and Fitch 2002; Fitch, Hauser, and Chomsky 2005); see also Hauser 1996. Biologists such as Salvador Luria (1973) and François Jacob (1982) have expressed a similar opinion, which goes back to the first cognitive revolution

- in the seventeenth and eighteenth centuries, where communication was often seen as a secondary role for language (its primary role being its relation to thought systems).
- 9 See Hauser 1996, and more explicitly Anderson 2004.
 - 10 Hauser 1996 is the most comprehensive survey I am familiar with.
 - 11 See Anderson 2004.
 - 12 See Hauser 1996, and, from the perspective of a linguist, see Anderson 2004.
 - 13 Taken from Anderson 2008.
 - 14 See Cheney and Seyfarth 1990; Hauser 1996; Anderson 2004.
 - 15 Marc Hauser, personal communication.
 - 16 On bees, see von Frisch 1967, 1974; Gould and Gould 1988; Hauser 1996; Anderson 2004.
 - 17 See Anderson 2004 and references therein.
 - 18 See Anderson 2004.
 - 19 See the survey in Eibl-Eibesfeldt 1970.
 - 20 This and other characteristics of bird songs are well documented in Anderson 2004.
 - 21 Bass, Gilland, and Baker 2008.
 - 22 See Dowling 2004.
 - 23 Chomsky 2007a, 2008.
 - 24 A point I stress in Boeckx to appear b.
 - 25 Hauser, Chomsky, and Fitch 2002; Fitch, Hauser, and Chomsky 2005.
 - 26 It is often said that Noam Chomsky dismisses talk of evolution of mental faculties (see, e.g., Newmeyer 1998). The written record indicates otherwise, as the collection of quotes (going back to works by Chomsky in the 1960s) in Otero 1990 and Jenkins 2000 attests.
 - 27 Hauser, Chomsky, and Fitch 2002.
 - 28 Gentner, Fenn, Margoliash and Nusbaum 2006.
 - 29 See, e.g., Anderson 2008; Liberman 2006; and Chomsky's reaction in Goudarzi 2006.
 - 30 Fitch and Hauser 2004.
 - 31 See, however, Hochmann, Azadpour, and Mehler 2008.
 - 32 For a reappraisal of "descent with modification" in the study of cognition, see Marcus 2004, 2006.
 - 33 On minute modifications leading to dramatic changes, see Carroll 2005 for an accessible treatment of this central tenet of the field of "evo-devo" (evolutionary-development biology), with ideas going back to Jacob and Monod 1961.
 - 34 For recent evidence that could be interpreted along these lines, see Fox and Hackl 2006 and especially Pietroski, Halberda, Lidz, and Hunter 2007.
 - 35 Piattelli-Palmarini and Uriagereka to appear.
 - 36 See Hauser 2006; Fitch 2005.
 - 37 As I have suggested in Boeckx to appear b.
 - 38 Lewontin 1998.
 - 39 As was recognized in the seventeenth and eighteenth centuries; see Schreyer 1985, 1989.

Chapter 12 Computational Organology

- 1 For good overviews of how complex the connection between genes and behavioral traits is, see Marcus 2004 and Rutter 2006.
- 2 Phillips 2004.

- 3 Chomsky has consistently stressed this obvious point; for a more recent, and in my view, quite compelling case, see Gallistel and King, 2009. See also Poeppel 2005, and Poeppel and Embick 2005, to whom I owe the title of this chapter.
- 4 Piattelli-Palmarini and Uriagereka 2005. For general reflections on FOXP2, see also Marcus and Fisher 2003; Fisher 2006; Piattelli-Palmarini and Uriagereka to appear; Berwick and Chomsky to appear. The text that follows draws heavily on Piattelli-Palmarini and Uriagereka to appear.
- 5 Piattelli-Palmarini and Uriagereka to appear.
- 6 See Pinker 1994; Jackendoff 1994.
- 7 Including a linguist (Myrna Gopnik; see Gopnik 1990).
- 8 I am here following Piattelli-Palmarini and Uriagereka's (to appear) assessment, who in turn rely on a comprehensive bibliography.
- 9 For a comprehensive description of the deficit, see Vargha-Khadem, Gadian, Copp, and Mishkin 2005.
- 10 For a review of the SLI literature, see van der Lely 2005.
- 11 See, e.g., Vargha-Khadem et al. 2005.
- 12 For a technical review, see Fisher, Lai, and Monaco 2003.
- 13 Lai et al. 2001.
- 14 Piattelli-Palmarini and Uriagereka to appear.
- 15 Piattelli-Palmarini and Uriagereka to appear.
- 16 For details and references to the technical literature, see Piattelli-Palmarini and Uriagereka to appear.
- 17 For a detailed attempt to identify other genes implicated, see Benítez Burraco 2009.
- 18 For a detailed critical overview, see Piattelli-Palmarini and Uriagereka to appear.
- 19 Enard et al. 2002.
- 20 Enard et al. 2002.
- 21 A point stressed by Piattelli-Palmarini and Uriagereka to appear, and Berwick and Chomsky to appear.
- 22 Piattelli-Palmarini and Uriagereka to appear.
- 23 Poeppel 2005; Poeppel and Embick 2005.
- 24 Hauser, Chomsky, and Fitch 2002.
- 25 Poeppel 2005.
- 26 As they have done in recent years; for review, see Boeckx 2006, Hornstein 2009.
- 27 A point stressed by Gallistel and King, 2009.
- 28 For a standard overview, see Kandel, Schwartz, and Jessell 2000. See also van Essen and Maunsell 1983; van Essen 1985; van Essen and Gallant 1994.
- 29 Piattelli-Palmarini and Uriagereka to appear.
- 30 I am here again following Piattelli-Palmarini and Uriagereka's (to appear) comprehensive review fairly closely.
- 31 Holy and Guo 2005; Shu et al. 2005.
- 32 Haesler et al. 2007.
- 33 Piattelli-Palmarini and Uriagereka to appear.
- 34 Poeppel 2005; Poeppel and Embick 2005.
- 35 Jacob 1977.
- 36 Marcus examines the implications of this fact for the emergence of our mental faculties in Marcus 2004, 2006.
- 37 Darwin 1859.

- 38 For striking evidence for neural recycling (from the visual domain to our reading ability), see Dehaene and Cohen 2007.
- 39 Rizzi 2004.
- 40 Marcus 2006.
- 41 For similar paradoxical situation involving language and numerical cognition, see Dehaene 1997.
- 42 Patel 2008.
- 43 As proposed by Marcus 2006.
- 44 Moro et al. 2001.
- 45 van der Lely 2005.
- 46 On this important difference, see Chomsky 2000, 2002.
- 47 Dehaene 2007.
- 48 Musso et al. 2003.
- 49 Smith and Tsimpli 1995.

Epilogue

- 1 Chomsky 1965.
- 2 Here I depart from Jackendoff's 2002 assessment that new foundational statements are called for.
- 3 Eldredge 2005.
- 4 For his most recent statement, see Jackendoff 2007.
- 5 That there may be a parallelism between our language sense and our moral sense was first mentioned in Rawls 1971, but this possibility wasn't systematically investigated until recently. I have found the works of Mikhail 2000, to appear; Dwyer 1999, 2006; and Hauser 2006 particularly insightful, and explicit about the connection with the study of the language faculty. Readers can find extensive evidence for the statements I made in the text in these works. For a popular review of our "moral instinct" (an obvious allusion to the "language instinct"), see Pinker 2008. As Mikhail 2000 points out, the idea that something like a computational procedure must be assumed in the moral domain goes back to David Hume.
- 6 Mikhail 2000.
- 7 Hauser 2006.
- 8 Dwyer 2006; Hauser 2006.
- 9 For the most comprehensive demonstration of this, see Hauser 2006.
- 10 Lerdahl and Jackendoff 1983. For an updated perspective, see Jackendoff and Lerdahl 2006, on which I draw in the text. For a very accessible overview of the range of issues that arise in the context of music cognition, see Jackendoff 1994.
- 11 Hauser, Chomsky and Fitch 2002.
- 12 Jackendoff and Lerdahl 2006.
- 13 Chomsky 1986.

Guide to Further Study

This study guide is intended to help readers (instructors and students) find ways to assimilate the material in this book and to explore further some issues I barely touched on. It is not to be followed rigidly, the way exercises are often presented in standard textbooks with a more technical bent than this one. Rather, I intend to provide a list of suggestions or pointers about ways to discuss the material from perspectives not explored directly or explicitly in the main text. Some of the suggestions were based on exercises I assigned in the various classes where I covered the material presented in this book, and proved successful and (by their own admissions) fun for students to explore.

Prologue and Chapter 1

❏ As a starting point, readers may want to reflect on ways the following questions could be answered:

- How do we acquire language?
- Why is it hard to learn a second language?
- Why is translation so difficult?
- Does language equal thought?
- Are sign languages real languages?
- Do animals have language?
- Can computers learn language?

These questions, taken from Part I of Napoli 2003, are all very complex and open-ended, hence often proving great sources for discussion in class, and can also be used as topics for short essays. Napoli provides answers that can be used to structure the discussion or act as positions to debate. I encourage the reader to revisit

these questions at various points in the book, and appreciate how some of the answers may change as they learn about some of the discoveries I report on.

❖ As the reader goes through Chapter 1, I would encourage her to find more examples of ambiguity, and maybe to find similar examples to those discussed in the text in other languages.

❖ Whenever I teach this material, I always try to give a list of examples (for example the binding cases in (6–11), one by one, to students, and let them come up with possible analyses which get confirmed or disconfirmed by the example. I have reproduced some of the examples here in the order I typically introduce them in class, and I hope instructors using this book will not ask students to read the chapter ahead of time, so as not to spoil for them the joy of discovering how intricate issues can emerge from simple sentences.

❖ Pullum's 2004 article on prescriptive and descriptive grammars can be a very good source of discussion. Readers are encouraged to identify additional discrepancies between prescriptive and descriptive rules.

Chapter 2

❖ Chomsky's 1959 review of Skinner's book is a must-read. I would in fact encourage students to read it twice: once to identify passages corresponding to the summary given in the text, and, in a second reading, to identify the passages that are inspired by ethology.

❖ I would also encourage readers to turn to chapter 1 of Chomsky's *Cartesian Linguistics*, "Creative Aspect of Language Use," and, if so inclined, read excerpts from Humboldt's 1836 long essay *On Language: On the Diversity of Human Language Construction and Its Influence on the Mental Development of the Human Species*.

❖ As an essay topic, I would consider assigning excerpts of the very beginning of Turing's proposal known as the Turing test (anthologized in Shieber 2004), and ask students to share their opinions.

❖ I would also recommend readers to read from Turing's 1936 paper on computation, and pay special attention to how much Turing seemed to address questions about what goes on in the brain. Petzold 2008 provides a book-length commentary on this seminal paper.

❖ Finally, it helps to read Piattelli-Palmarini 1987, 1989; and Jerne 1985 to drive home the point about selection in learning.

Chapter 3

❖ Plato's *Meno* is a must-read (at least the part when Socrates raises questions to the slave boy). Instructors should encourage students to replace Plato's original discussion about innate mathematical knowledge with examples from other cognitive domains (anticipating the Epilogue).

❏ I would also recommend reading Lenneberg's shorter essays such as Lenneberg 1960, 1964, 1966, and relating the themes discussed there to the discussion of ethology in Chapter 2, and Chomsky's 1959 review. Drawing parallels between Lenneberg's works and Chomsky's review is a very valuable exercise.

❏ I also recommend students to browse through child speech in the CHILDES database (childes.psy.cmu.edu), select a few sections of one or two files (CHILDES is very large) and identify as many interesting properties in the child's production as they can; this is an exercise to be repeated after Chapter 6 has been covered.

❏ For an accessible, engaging, and still very relevant debate on learning, see Piattelli-Palmarini 1980.

❏ Comparing utterances from pidgins and creoles provides the best kind of illustration of many of the points made in the text. Consider the case of Hawaiian creole (from Bickerton's work; also discussed in Jackendoff 1994):

- a. Pidgin: No, the men, ah-pau [= finished] work-they go, make garden. Plant this, ah, cabbage, like that. Plant potato, like that. And then-all that one-all right, sit down. Make lilly bit story.
- b. Creole: When work pau [= is finished] da guys they stay go make garden for plant potato an' cabbage an' after little while they go sit talk story

(Reader should focus on the way information is integrated structurally.)

❏ I encourage readers to collect examples in other languages of pairs of minimally different words that depend on a sound contrast that is missing in their native language.

Chapter 4

❏ In my experience attempting to read Chomsky's *Syntactic Structures* without much background can be a frustrating experience. I recommend using Lasnik 2000 as a point of entry.

❏ Lasnik 2000 contains excellent exercises, including some pertaining to the construction of finite-state machines and rewrite-rule systems. For example, readers may want to consider constructing a finite-state machine for something like "My friend cried (and cried, and cried, and . . .)." Readers may also attempt to construct a rewrite rule system (as compact as possible) for what is sometimes called a mirror-image language ("Martian") like:

aa
bb
abba
baab
abbbba
...

Instructors may also want to use a different fragment of English from the one used in the text and ask students to construct a rewrite rule system for it, along with tree diagrams.

❖ All the properties of grammar discussed in the text (locality, endocentricity, etc.) form the standard subject matter of introductory syntax textbooks, such as Carnie 2007, Adger 2003, and Haegeman 2006, where readers can find a wealth of exercises. Because this book is not meant to be a full-blown introduction to technical aspects of the field of linguistics (it is only meant to prepare students to appreciate the need to study such aspects), I recommend against using Chapter 4 as a way of turning the material into a mini-introduction to syntax. It is important to let students reflect on some of the phenomena discussed in this chapter, and let them come up with other examples that may illustrate the same points. In the context of locality, I have always found it useful to let students reflect on more cases of what Ross 1967 called islands, and try to find ways to explain why they are unacceptable (students should resist the temptation to say that such sentences don't make sense; it is easy to imagine contexts where asking such questions would make a lot of sense). For example:

- a. Who did pictures of annoy Bill? (compare: Who did Bill see pictures of?)
- b. Who did John say that left? (compare: Who did John say left?)
- c. Who did Sue arrive after Bill kissed? (compare: Who did Bill kiss before Sue arrived?)
- d. Whose did you see book? (compare: Whose book did you see?)

Chapter 5

The ideal companion to this chapter is Baker 2001.

❖ A first good exercise would be to try to pick one of the parameters not discussed in this book, but discussed in Baker, and try to explain it in your own terms. I would insist on the notion of “dependency” among parameters; try to indentify as many cases of dependency as possible in Baker 2001.

❖ I would also encourage readers to try to describe some of the differences between English and another language they know in terms of parameters (perhaps proposing new parameters along the way).

❖ In the main text I discuss head-final languages like Japanese, but I do not discuss head-initial languages (Verb-Subject-Object languages). Try to describe them in your own terms, then turn to Baker 2001, who discusses relevant examples.

❖ Turn passages of English into “pseudo-Japanese”: keep the English words but pretend English is head-final. For example, what would the following example look like:

Girls that were proud of their countries kissed men next to them.

❏ Baker 2001 discusses the serial verb parameter, which allows languages to put more than one verb in the same sentence/clause without using any overt conjunction. English is not one of the languages that allow this, but if it did, it would allow sentences like “John went cook the meal.”

There is, however, a small corner of English grammar where serial verbs are possible. Consider:

Go get me the book.

Interestingly, Baker notes in his discussion of serial verbs that languages that allow them do not mark tense directly on verbs (they use a separate little word for that, as English sometimes does: *John did not leave*). Now consider the fact that although English allows *go get me the book* it does not allow *John went got me the book* (but it allows *John will go get me the book*). Discuss the implications of these examples for Baker’s classification of English as saying “no” to serial verbs.

Chapter 6

❏ For readers who want to find out more about language acquisition, I recommend Guasti 2002. O’Grady 2005 also provides a wealth of interesting data and studies.

❏ The best exercise to appreciate some of the findings reported on in this chapter is to confront data from language acquisition and language change. For example, take an excerpt from older varieties of English (say, an excerpt from one of Shakespeare’s plays), and try to describe in terms of parameters some of the changes that took place to arrive at Modern English. Also take a look at transcripts from the CHILDES database, and identify ways in which Child English departs from adult English, and again, try to describe the difference in terms of parameters. To appreciate the difficulty of gathering experimental data with children, see Crain and Thornton 1998.

Chapter 7

❏ This chapter, more than any other, has a distinct philosophical orientation. As such this is the best place to let students discuss and debate the various opinions on the table, and maybe tackle some of Chomsky’s own writings on the matter. I recommend assigning Stainton 2006, and then turning to Chomsky 2000.

❏ I would also encourage students/readers to collect more data on negative polarity items, trying out the examples discussed in the text with different determiners/quantifiers (*all, both, many, most, each*, etc.)

❏ Given the emphasis laid on negative facts in this chapter, I would encourage students/readers to collect more instances of ambiguous sentences and identify the interpretations that are not available, as well as any other negative facts they can

think of. Since this is the first chapter to touch on concepts, I cannot fail to urge readers to read Fodor's claim that all (basic) concepts we have are innate (see Fodor 1975, see also his contributions to Piattelli-Palmarini 1980). I do not discuss Fodor's claim in the text because he has done a better job at discussing it and its implications than I would ever be able to do. I strongly recommend his chapter on innateness in Fodor 2008.

Chapter 8

❖ Fodor's 1983 *The Modularity of Mind* is a classic, and deserves to be read. I would also recommend Hauser's 2000 *Wild Minds*. The latter is divided into three parts: what is common among many animals (core knowledge systems) (Part I), what is less clearly shared (including chapters on deception, imitation and self-recognition) (Part II), and what is unique to humans (Part III). Part II is the most controversial section of the book, and given that it was published in 2000, it would be a good exercise to try to find out if we have learned more in these domains more recently. This chapter is also an ideal starting point for students to appreciate all the subtleties involved in experimental designs. I would assign a few of the experiment-oriented papers by Spelke and her collaborators and focus on how experiments must be carefully constructed to arrive at the conclusions reported on in this chapter.

Chapter 9

❖ I urge readers to read Fromkin's 1973 classic on speech errors and what they reveal about linguistic knowledge, and how it is put to use. For an excellent survey of issues in sentence processing, I recommend Townsend and Bever 2001.

❖ To appreciate the subtlety of contrasts to be dealt with in any study relying on speakers' intuition, readers should judge (and ask other native speakers to judge) the following and discuss any pattern they may recognize.

- a. Who did John say that Mary saw?
- b. Who did John say Mary saw?
- c. Who did John say that saw Mary?
- d. Who did John say saw Mary?
- e. Who did John wonder whether Mary saw?
- f. Who did John wonder whether saw Mary?
- g. Which girl did John wonder whether Bill saw?
- h. Which of the two girls did John wonder whether Bill saw?
- i. What did John ask how to cook?
- j. Who asked who bought what?
- k. Who asked what who bought?

Chapter 10

❏ I urge readers to read the historical pieces by Broca and the other fathers of the classic model collected in Grodzinsky and Amunts 2006 to get a feeling of what they were after and how modern many of them read.

❏ For a delightful discussion of the intricacies of running neurolinguistic experiments and how they might bear on the big questions, see Moro 2008, especially chapter 2, which would be an ideal complement to this chapter.

❏ I also encourage readers to turn to the neurolinguistic literature (there is a profusion of journals in this area) and try to identify the explicit or implicit assumptions that form part of the new phrenology.

❏ Finally, a closer look at aphasia studies may be a good way to branch off. Below are a few samples of speech from a Broca-aphasic and a Wernicke-aphasic.

Broca's aphasia:

"Yes . . . Monday . . . Dad, and Dad . . . hospital, and . . . Wednesday, Wednesday, nine o'clock and . . . Thursday, ten o'clock . . . doctors, two, two . . . doctors and . . . teeth, yah. And a doctor . . . girl, and gums, and I."

"Me . . . build-ing . . . chairs, no, no cab-in-ets. One, saw . . . then, cutting wood . . . working . . ."

Wernicke's aphasia:

EXAMINER: What kind of work have you done?

PATIENT: We, the kids, all of us, and I, we were working for a long time in the . . . you know . . . it's the kind of space, I mean place rear to the spedawn . . .

EXAMINER: Excuse me, but I wanted to know what work you have been doing.

PATIENT: If you had said that, we had said that, poomer, near the fortunate, por-punate, tamppoo, all around the fourth of martz. Oh, I get all confused.

Chapter 11

❏ I cannot fail to recommend the reader of this chapter on evolutionary questions to turn to Darwin's own writings. Many passages of *The Descent of Man* (1871) are still extremely relevant, and contain a wealth of insights that still await examination in a modern context. I also urge the reader to consult Anderson 2004, the ideal complement to this chapter. For readers who want more details about systems of communication in other animals, Hauser 1996 is a must-read.

❏ On the limits of adaptationism, I recommend Lewontin 1998.

❏ As a useful exercise, I would encourage readers to read, compare, and debate the positions in Hauser, Chomsky, and Fitch 2002 and Pinker and Jackendoff 2005, as well as the reply and rejoinder in Fitch, Hauser, and Chomsky 2005 and Jackendoff and Pinker 2005.

Chapter 12

❖ This chapter mentions a few terms from genetics that readers may want to examine more closely by turning to genetics tutorials. There is a profusion of such tutorials online, but I also recommend the following two books, written from a cognitive perspective: Marcus 2004 and Rutter 2006. I also urge the readers to read Lewontin's delightful books on the complexity of gene–environment interactions (Lewontin 2000a, 2000b). They should serve as an ideal antidote against the genocentric view to be found in so many popular books about human nature.

Epilogue

❖ In addition to summarizing some of the main points in this book, this final chapter refers to recent work on music cognition and morality. Readers with a musical bent and some knowledge of music terminology should turn to Lerdahl and Jackendoff's 1983 classic study. Works on the moral sense are likely to prove very controversial. The ideal starting point would be to go the moral sense test (<http://moral.wjh.harvard.edu/>) concocted by Marc Hauser and collaborators and read the discussion of some of the findings in Hauser 2006. Pinker 2008 could also serve as an excellent position paper to start the discussion. Because so much work remains to be done at the basic level in this domain, readers are encouraged to think of ways of formulating principles and parameters along the lines pursued in this book for language.

❖ This final chapter is also the idea point to wrap up with an overview exercise I have often tried in class. For an instructor the biggest challenge is not so much what questions to put on the exam (i.e., figure out what one wants students to remember tomorrow and, no doubt, forget immediately afterwards), but rather to be clear about what concepts and issues one would like students to remember, say, a year from taking the class. It's the same question an author of a book faces: what are the main points to take away from the book?

It would help students/readers zero in on the essentials of this book if they could attempt something like a re-writing of Plato's *Meno* dialogue, imagining a friend who hasn't read this book as interlocutor, and trying to incorporate the key points and a few salient illustrations that they have derived from this book. I will resist the temptation to list all the key concepts introduced in each chapter, and let the readers identify them for themselves. For those tempted by this overview exercise, I urge them not to replicate the exact same order in which the key ideas were introduced; come up with a different "table of contents" and also different examples to illustrate some of the main points. Remember that I had a hard time selecting the material for this book due to the profusion of excellent works that are more and more easily accessible online. This is an ideal chance to be creative.

References

- Adger, D. 2003. *Core syntax: a minimalist approach*. Oxford: Oxford University Press.
- Anderson, S. R. 2004. *Doctor Dolittle's delusion*. New Haven, CT: Yale University Press.
- Anderson, S. R. 2008. The logical structure of linguistic theory. *Language* 84: 795–814.
- Anderson, S. R. & Lightfoot, D. 2002. *The language organ*. Cambridge: Cambridge University Press.
- Anwander, A., Tittgemeyer, M., von Cramon, D., Friederici, A., & Knösche, T. 2007. Connectivity-based parcellation of Broca's area. *Cerebral Cortex* 17: 816–25.
- Arthur, W. 2004. *Biased embryos and evolution*. Cambridge: Cambridge University Press.
- Aslin, R. N., Pisoni, D. B., Hennessy, B. I., & Perey, A. J. 1981. Discrimination of voice onset time by human infants: new findings and implications for the effects of early experience. *Child Development* 52: 1135–45.
- Baillargeon, R. 1987. Object permanence in 3.5- and 4.5-month-old infants. *Developmental Psychology* 23: 655–64.
- Baillargeon, R. 2008. Innate ideas revisited: for a principle of persistence in infants' physical reasoning. *Perspectives on Psychological Science* 3: 2–12.
- Baker, M. 1996. *The polysynthesis parameter*. Oxford: Oxford University Press.
- Baker, M. 2001. *The atoms of language: the mind's hidden rules of grammar*. New York: Basic Books.
- Baker, M. 2003. Language differences and language design. *Trends in Cognitive Sciences* 7: 349–53.
- Baker, M. 2005. Mapping the terrain of language acquisition. *Language Learning and Development* 1: 93–129.
- Barkow, J. H., Cosmides, L., & Tooby, J., eds. 1992. *The adapted mind: evolutionary psychology and the generation of culture*. Oxford: Oxford University Press.
- Barwise, J. & Cooper, R. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 4: 159–219.
- Bass, A. H., Gilland, E. H., & Baker, R. 2008. Evolutionary origins for social vocalization in a vertebrate hindbrain-spinal compartment. *Science* 321: 417–21.
- Ben Shalom, D. & Poeppel, D. 2008. Functional anatomic models of language: assembling the pieces. *Neuroscientist* 14: 119–27.

- Benítez Burraco, A. 2009. *Genes y lenguaje: Aspectos ontogenéticos, filogenéticos, y cognitivos*. Barcelona: Editorial Reverté.
- Benzaquén, A. S. 2006. *Encounters with wild children: temptation and disappointment in the study of human nature*. Montréal: McGill-Queen's University Press.
- Bermúdez, J. L. 2003. *Thinking without words*. Oxford: Oxford University Press.
- Berwick, R. C. 1985. *The acquisition of syntactic knowledge*. Cambridge, MA: MIT Press.
- Berwick, R. C. & Chomsky, N. To appear. The Biolinguistic program: the current state of its development. In A.-M. Di Sciullo & C. Boeckx, eds., *The biolinguistic enterprise*. Oxford: Oxford University Press.
- Berwick, R. C. & Niyogi, P. 1996. Learning from triggers. *Linguistic Inquiry* 27: 605–22.
- Bever, T. G. 1970. The cognitive basis for linguistic structures. In R. Hayes, ed., *Cognition and language development* (pp. 279–360). New York: Wiley.
- Bickerton, D. 1981. *Roots of language*. Ann Arbor: Karoma.
- Bickerton, D. 1984. The language bioprogram hypothesis. *Behavioral and Brain Sciences* 7: 173–88.
- Biegler, R. & Morris, R. G. 1993. Landmark stability is a prerequisite for spatial but not discrimination learning. *Nature* 361: 631–3.
- Biegler, R. & Morris, R. G. 1996. Landmark stability: further studies pointing to a role in spatial learning. *Quarterly Journal of Experimental Psychology B* 49: 307–45.
- Bilgrami, A. 2002. Chomsky and philosophy. *Mind and Language* 17: 290–302.
- Bloom, P. 2000. *How children learn the meanings of words*. Cambridge, MA: MIT Press.
- Bloom, P. 2004. *Descartes' baby: how the science of child development explains what makes us human*. New York: Basic Books.
- Boden, M. 2006. *Mind as machine: a history of cognitive science* (2 vols.). Oxford: Oxford University Press.
- Boeckx, C. 2006. *Linguistic minimalism: origins, concepts, methods, and aims*. Oxford: Oxford University Press.
- Boeckx, C. 2008. *Bare syntax*. Oxford: Oxford: Oxford University Press.
- Boeckx, C. 2009. The nature of Merge: consequences for language, mind, and biology. In M. Piattelli-Palmarini, J. Uriagereka, & P. Salaburu, eds., *Of minds and language: the Basque Country encounter with Noam Chomsky* (pp. 44–57). Oxford: Oxford University Press.
- Boeckx, C. To appear a. Approaching parameters from below. In A.-M. Di Sciullo & C. Boeckx, eds., *The biolinguistic enterprise*. Oxford: Oxford University Press.
- Boeckx, C. To appear b. Some reflections on Darwin's problem in the context of Cartesian Biolinguistics. In A.-M. di Sciullo and C. Boeckx, eds., *The biolinguistic enterprise*. Oxford: Oxford University Press.
- Boeckx, C. & Grohmann, K. K., eds. 2003. *Multiple wh-fronting*. Amsterdam: John Benjamins.
- Boeckx, C. & Grohmann, K. K. 2007. The *Biolinguistics* manifesto. *Biolinguistics* 1: 1–8.
- Boeckx, C. & Hornstein, N. 2007. Les différents objectifs de la linguistique théorique. In J. Frank & J. Bricmont, eds., *Cahier Chomsky* (pp. 61–77). Paris: L'Herne.
- Boeckx, C. & Piattelli-Palmarini, M. 2007. Linguistics in cognitive science: the state of the art amended. *The Linguistic Review* 24: 403–16.
- Borer, H. & Wexler, K. 1987. The maturation of syntax. In T. Roeper & E. Williams, eds., *Parameter setting* (pp. 123–72). Dordrecht: Reidel.
- Borer, H. & Wexler, K. 1992. Bi-unique relations and the maturation of grammatical principles. *Natural Language and Linguistic Theory* 10: 147–89.

- Boyer, P. 2003. *Religion explained*. New York: Basic Books.
- Bresnan, J. 2007. Is syntactic knowledge probabilistic? Experiments with the English dative alternation. In S. Featherston & W. Sternefeld, eds., *Roots: linguistics in search of its evidential base* (pp. 75–96). Berlin: Mouton de Gruyter.
- Broca, P. 1861. Loss of speech, chronic softening, and partial destruction of the anterior left lobe of the brain. *Bulletin de la Société Anthropologique* 2: 235–8.
- Brodmann, K. 1908. *Die vergleichende Lokalisationslehre der Großhirnrinde*. Leipzig: Barth.
- Brook, A. 2007. *The prehistory of cognitive science*. Basingstoke, UK: Palgrave Macmillan.
- Burge, T. 1979. Individualism and the mental. *Midwest Studies in Philosophy* 4: 73–121.
- Büring, D. 2005. *Binding theory*. Cambridge: Cambridge University Press.
- Butterworth, B. 1999. *What counts: how every brain is hardwired for math*. New York: The Free Press.
- Campbell, L. & Poser, W. J. 2008. *Language classification: history and method*. Cambridge: Cambridge University Press.
- Caramazza, A. & Zurif, E. B. 1976. Dissociations of algorithmic and heuristic processes in sentence comprehension: evidence from aphasia. *Brain and Language* 3: 572–82.
- Carey, S. 1986. Cognitive science and science education. *American Psychologist* 41: 1123–30.
- Carey, S. & Spelke, E. S. 1994. Domain-specific knowledge and conceptual change. In L. Hirschfeld & S. Gelman, eds., *Mapping the mind: domain specificity in cognition and culture* (pp. 169–200). Cambridge: Cambridge University Press.
- Carey, S. & Spelke, E. S. 1996. Science and core knowledge. *Philosophy of Science* 63: 515–33.
- Carey, S., Honda, M., Jay, E., O’Neil, W., & Unger, C. 1986. What junior high school students do, can, and should know about the nature of science. Educational technology center technical report 86–11. Harvard Graduate School of Education.
- Carnie, A. 2007. *Syntax*, 2nd edn. Oxford: Blackwell.
- Carroll, L. 1872. *Through the looking-glass, and what Alice found there*. London: Macmillan.
- Carroll, S. B. 2005. *Endless forms most beautiful: the new science of evo devo and the making of the animal kingdom*. New York: W. W. Norton.
- Carruthers, P. 2002. The cognitive functions of language. *Behavioral and Brain Sciences* 25: 657–73.
- Carruthers, P. 2006. *The architecture of the mind: massive modularity and the flexibility of thought*. Oxford: Oxford University Press.
- Changeux, J.-P., Courrège, P., & Danchin, A. 1973. A theory of epigenesis of neuronal networks by selective stabilization of synapses. *Proceedings of the National Academy of Sciences* 70: 2974–8.
- Chater, N., Tenenbaum, J. B., & Yuille, A. 2006. Probabilistic models of cognition: conceptual foundations. *Trends in Cognitive Sciences* 10: 287–91.
- Cheney, D. & Seyfarth, R. 1990. *How monkeys see the world*. Chicago: University of Chicago Press.
- Cheng, K. 1986. A purely geometric module in the rat’s spatial representation. *Cognition* 23: 149–78.
- Cherbury, H. 1624. *De veritate, prout distinguitur a revelatione, a verisimili, a possibili, et a falso*. Paris.
- Chierchia, G. & McConnell-Ginet, S. 1990. *Meaning and grammar: an introduction to semantics*. Cambridge, MA: MIT Press.

- Chomsky, C., Honda, M., O'Neil, W., & Unger, C. 1985. Doing science: constructing scientific theories as an introduction to scientific method. Educational technology center technical report 85-23. Harvard Graduate School of Education.
- Chomsky, N. 1955. The logical structure of linguistic theory. Ms., Harvard/MIT. (Published in part, 1975, New York: Plenum.)
- Chomsky, N. 1956. Three models for the description of language. *IRE Transactions on Information Theory*, vol. IT-2, *Proceedings of the Symposium on Information Theory* 113–24.
- Chomsky, N. 1957. *Syntactic structures*. The Hague: Mouton.
- Chomsky, N. 1959. Review of B. F. Skinner's *Verbal behavior*. *Language* 35: 26–58.
- Chomsky, N. 1964. *Current issues in linguistic theory*. The Hague: Mouton.
- Chomsky, N. 1965. *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.
- Chomsky, N. 1966. *Cartesian linguistics*. New York: Harper & Row.
- Chomsky, N. 1968. *Language and mind*. New York: Harcourt, Brace, & World.
- Chomsky, N. 1970. Remarks on nominalization. In R. Jacobs & P. Rosenbaum, eds., *Readings in English transformational grammar* (pp. 184–221). Waltham, MA: Blaisdell Publishing.
- Chomsky, N. 1971. *Problems of knowledge and freedom*. New York: New Press.
- Chomsky, N. 1973. Conditions on transformations. In S. Anderson & P. Kiparsky, eds., *A festschrift for Morris Halle* (pp. 232–86). New York: Holt, Rinehart, & Winston.
- Chomsky, N. 1975. *Reflections on language*. New York: Pantheon.
- Chomsky, N. 1976. Conditions on rules of grammar. *Linguistic Analysis* 2: 303–51.
- Chomsky, N. 1977. On wh-movement. In P. Culicover, T. Wasow, & A. Akmajian, eds., *Formal syntax* (pp. 71–13). New York: Academic Press.
- Chomsky, N. 1980. *Rules and representations*. New York: Columbia University Press.
- Chomsky, N. 1981. *Lectures on government and binding*. Dordrecht: Foris.
- Chomsky, N. 1986. *Knowledge of language*. New York: Praeger.
- Chomsky, N. 1988. *Language and problems of knowledge*. Cambridge, MA: MIT Press.
- Chomsky, N. 1991. Some notes on economy of derivation and representation. In R. Freidin, ed., *Principles and parameters in comparative grammar* (pp. 417–54). Cambridge, MA: MIT Press. (Reprinted in N. Chomsky, 1995, *The minimalist program* (pp. 129–66). Cambridge, MA: MIT Press.)
- Chomsky, N. 1993. A minimalist program for linguistic theory. In K. Hale & S. J. Keyser, eds., *The view from Building 20* (pp. 1–52). Cambridge, MA: MIT Press. (Reprinted in N. Chomsky, 1995, *The minimalist program* (pp. 167–217). Cambridge, MA: MIT Press.)
- Chomsky, N. 1995. *The minimalist program*. Cambridge, MA: MIT Press.
- Chomsky, N. 2000. *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.
- Chomsky, N. 2002. *On nature and language*. Cambridge: Cambridge University Press.
- Chomsky, N. 2004. Language and mind: current thoughts on ancient problems. In L. Jenkins, ed., *Variation and universals in biolinguistics* (pp. 379–406). Amsterdam: Elsevier.
- Chomsky, N. 2005. Three factors in language design. *Linguistic Inquiry* 36: 1–22.
- Chomsky, N. 2007a. Approaching UG from below. In U. Sauerland & H. M. Gärtner, eds., *Interfaces + recursion = language?* (pp. 1–29). Berlin: Mouton de Gruyter.
- Chomsky, N. 2007b. Book review. Symposium on Margaret Boden: *Mind as machine*. *Artificial Intelligence* 1094–1103.
- Chomsky, N. 2008. On phases. In R. Freidin, C. Otero, & M.-L. Zubizarreta, eds., *Foundational issues in linguistic theory* (pp. 133–66). Cambridge, MA: MIT Press.

- Chomsky, N. In press. The mysteries of nature: how deeply hidden? *Journal of Philosophy*.
- Chomsky, N. & Lasnik, H. 1993. The theory of principles and parameters. In J. Jacobs, A. von Stechow, W. Sternefeld, & T. Vennemann, eds., *Syntax: an international handbook of contemporary research* (pp. 506–69). Berlin: Mouton de Gruyter.
- Chomsky, N. & Miller, G. A. 1963. Introduction to the formal analysis of natural languages. In R. D. Luce, R. R. Bush, & E. Galanter, eds., *Handbook of mathematical psychology* (vol. 2, pp. 269–321). New York: Wiley.
- Christiansen, M. H. & Kirby, S., eds. 2003. *Language evolution*. Oxford: Oxford University Press.
- Christianson, K., Hollingworth, A., Halliwell, J., & Ferreira, F. 2001. Thematic roles assigned along the garden path linger. *Cognitive Psychology* 42: 368–407.
- Christophe, A., Guasti, M. T., Nespor, M., Dupoux, E., & van Ooyen, B. 1997. Reflections on phonological bootstrapping: its role for lexical and syntactic acquisition. *Language and Cognitive Processes* 12: 585–612.
- Churchland, P. S. 1994. Presidential address of the APA Pacific Division.
- Cinque, G. 1990. *Types of A-bar dependencies*. Cambridge, MA: MIT Press.
- Cinque, G. 1999. *Adverbs and functional heads*. Oxford: Oxford University Press.
- Cinque, G. & Kayne, R., eds. 2005. *The Oxford handbook of comparative syntax*. Oxford: Oxford University Press.
- Clark, R. & Roberts, I. 1993. A computational approach to language learnability and language change. *Linguistic Inquiry* 24: 299–345.
- Cottingham, J. 1978. A brute to the brutes? Descartes' treatment of animals. *Philosophy* 53: 551–61.
- Crain, S. 1991. Language acquisition in the absence of experience. *Behavioral and Brain Sciences* 4: 597–650.
- Crain, S. & Lillo-Martin, D. 1999. *An introduction to linguistic theory and language acquisition*. Oxford: Blackwell.
- Crain, S. & Pietroski, P. 2001. Nature, nurture, and universal grammar. *Linguistics and Philosophy*, 24: 139–86.
- Crain, S. & P. Pietroski. 2002. Why language acquisition is a snap. *The Linguistic Review* 19: 163–83.
- Crain, S. & Thornton, R. 1991. Recharting the course of language acquisition: studies in elicited production. In N. Krasnegor, D. Rumbaugh, R. Schiefelbusch, & M. Studdert-Kennedy, eds., *Biobehavioral foundations of language development* (pp. 321–37). Hillsdale, NJ: Lawrence Erlbaum.
- Crain, S. & Thornton, R. 1996. Why child grammar is flawed. In C. Koster & F. Wijnen, eds., *Proceedings of GALA 1995*. Groningen: Center for Language & Cognition.
- Crain, S. & Thornton, R. 1998. *Investigations in universal grammar: a guide to experiments in the acquisition of syntax and semantics*. Cambridge, MA: MIT Press.
- Crain, S., Goro, T., & Thornton, R. 2006. Language acquisition is language change. *Journal of Psycholinguistic Research* 35: 31–49.
- Crick, F. 1994. *The astonishing hypothesis: the scientific search for the soul*. New York: Touchstone.
- Cudworth, R. 1688 [1996]. *A treatise concerning eternal and immutable reality: with a treatise of freewill*, ed. S. Hutton, Cambridge: Cambridge University Press.
- Curtiss, S., ed. 1977. *Genie: psycholinguistic study of a modern-day "wild child."* London: Academic Press.

- Curtiss, S., Fromkin, V., Krashen, S., Rigler, D., & Rigler, M. 1974. The linguistic development of Genie. *Language* 50: 528–54.
- Damasio, A. 1994. *Descartes' error*. New York: Putnam.
- Damasio, A. 2003. *Looking for Spinoza: joy, sorrow, and the feeling brain*. New York: Harcourt.
- Darwin, C. 1859. *On the origin of species by means of natural selection*. London: John Murray.
- Darwin, C. 1871. *The descent of man, and selection in relation to sex*. London: John Murray.
- Dascal, M. 2007. Hobbes's challenge. In A. Brook, ed., *The prehistory of cognitive science* (pp. 67–96). Basingstoke, UK: Palgrave Macmillan.
- Davidson, D. 1967a. Truth and meaning. *Synthese* 17: 304–23.
- Davidson, D. 1967b. The logical form of action sentences. In N. Rescher, ed., *The logic of decision and action* (pp. 81–95). Pittsburgh: University of Pittsburgh Press.
- Davis, J. 1997. *Mapping the mind: the secrets of the human brain and how it works*. Secaucus, NJ: Carol Publishing Group.
- Davis, M. 2000. *Engines of logic*. New York: Norton.
- Davis, M., ed. 2004. *The undecidable: basic papers on undecidable propositions, unsolvable problems, and computable functions*. New York: Dover.
- Dawkins, R. 1976. *The selfish gene*. Oxford: Oxford University Press.
- Dawkins, R. 1996. *The blind watchmaker*. New York: Norton.
- Dawkins, R. 1998. *Unweaving the rainbow: science, delusion, and the appetite for wonder*. Boston: Mariner Books.
- DeCasper, A. J. & Fifer, W. P. 1980. Of human bonding: newborns prefer their mothers' voices. *Science* 208: 1174–6.
- DeCasper, A. J. & Spence, M. J. 1986. Prenatal maternal speech influences newborns' perception of speech sounds. *Infant Behavior and Development* 9: 133–50.
- DeGraff, M., ed. 1999. *Language creation and language change: creolization, diachrony, and development*. Cambridge, MA: MIT Press.
- Dehaene, S. 1997. *The number sense*. Oxford: Oxford University Press.
- Dehaene, S. 2007. A few steps towards a science of mental life. *Mind, Brain, and Education* 1: 28–47.
- Dehaene, S. & Cohen, L. 2007. Cultural recycling of cortical maps. *Neuron* 56: 384–98.
- Dennett, D. 1995. *Darwin's dangerous idea: evolution and the meanings of life*. New York: Simon & Schuster.
- Descartes, R. 1641. *Meditationes de prima philosophia*. Leiden.
- Descartes, R. 1985. *The philosophical writings of Descartes* (2 vols.), trans. J. Cottingham, R. Stoothoff, and D. Murdoch. Cambridge: Cambridge University Press.
- Dijksterhuis, E. J. 1961. *The mechanization of the world picture*. Oxford: Clarendon Press.
- Dowling, J. E. 1998. *Creating mind: how the brain works*. New York: W. W. Norton.
- Dowling, J. E. 2004. *The great brain debate: is it nature or nurture?* Washington, DC: Joseph Henry Press.
- Dresher, B. E. 1999. Charting the learning path: cues to parameter setting. *Linguistic Inquiry* 30: 27–67.
- Dresher, B. E. 2003. Meno's paradox and the acquisition of grammar. In S. Ploch, ed., *Living on the edge: 28 papers in honour of Jonathan Kaye* (pp. 7–27), *Studies in Generative Grammar* 62. Berlin: Mouton de Gruyter.
- Dresher, B. E. & Kaye, J. D. 1990. A computational learning model for metrical phonology. *Cognition* 34: 137–95.

- Dwyer, S. 1999. Moral competence. In K. Murasugi & R. Stainton, eds., *Philosophy and linguistics* (pp. 169–90). Boulder, CO: Westview Press.
- Dwyer, S. 2006. How good is the linguistic analogy? In P. Carruthers, S. Laurence, & S. Stich, eds., *The innate mind: culture and cognition* (pp. 237–56). New York: Oxford University Press.
- Edelman, G. 1978. Group selection and phasic re-entrant signalling: a theory of higher brain function. In V. Mountcastle, ed., *The mindful brain* (pp. 51–100). Cambridge, MA: MIT Press.
- Edelman, G. 1989. *Neural Darwinism*. Oxford: Oxford University Press.
- Eibl-Eibesfeldt, I. 1970. *Ethology: the biology of behavior*. New York: Holt, Rinehart, & Winston.
- Eibl-Eibesfeldt, I. 1989. *Human ethology*. New York: Aldine de Gruyter.
- Eimas, P., Siqueland, E., Jusczyk, P., & Vigorito, J. 1971. Speech perception in infants. *Science* 171: 303–6.
- Eisenberg, R. B. 1975. *Auditory competence in early life*. Baltimore: University Park Press.
- Eldredge, N. 2005. Patterns. *Natural History* 114: 80.
- Elman, J., Bates, E., Johnson, M., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. 1997. *Rethinking innateness: a connectionist perspective on development*. Cambridge, MA: MIT Press.
- Embick, D. & Poeppel, D. 2006. Mapping syntax using imaging: prospects and problems for the study of neurolinguistic computation. In K. Brown, ed., *Encyclopedia of language and linguistics* (2nd edn, vol. 7, pp. 484–6). Oxford: Elsevier.
- Emonds, J. 1978. The verbal complex V'–V in French. *Linguistic Inquiry* 9: 151–75.
- Enard, W., Przeworski, M., Fisher, S. E., Lai, C. S. L., Wiebe, V., Kitano, T., Monaco, A. P., & Paabo, S. 2002. Molecular evolution of FOXP2, a gene involved in speech and language. *Nature* 418: 869–72.
- Essen, D. C. van 1985. Functional organization of primate visual cortex. In A. Peters and E. G. Jones, eds., *Cerebral cortex* (vol. 3, pp. 259–329). New York: Plenum.
- Essen, D. C. van & Gallant, J. L. 1994. Neural mechanisms of form and motion processing in the primate visual system. *Neuron* 13: 1–10.
- Essen, D. C. van & Maunsell, J. H. R. 1983. Hierarchical organization and functional streams in the visual cortex. *Trends in Neuroscience* 6: 370–5.
- Fabb, N. 1985. Linguistics for ten-year-olds. In D. Archangeli, A. Barss, and R. Sproat, eds., *Papers in theoretical and applied linguistics* (MIT Working Papers in Linguistics, 6, pp. 45–61). Cambridge, MA: MITWPL.
- Fauconnier, G. 1975. Pragmatic scales and logical structure. *Linguistic Inquiry* 6: 353–75.
- Feigenson, L. 2007. The equality of quantity. *Trends in Cognitive Sciences* 11: 185–7.
- Feigenson, L. & Carey, S. 2003. Tracking individuals via object-files: evidence from infants' manual search. *Developmental Science* 6: 568–84.
- Feigenson, L. & Carey, S. 2005. On the limits of infants' quantification of small object arrays. *Cognition* 97: 295–313.
- Feigenson, L., Carey, S., & Hauser, M. 2002. The representations underlying infants' choice of more: object-files versus analog magnitudes. *Psychological Science* 13: 150–6.
- Feigenson, L., Dehaene, S., & Spelke, E. S. 2004. Core systems of number. *Trends in Cognitive Sciences* 8: 307–14.
- Fellerman, D. J. & Essen, D. C. van 1991. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex* 1: 1–47.
- Ferreira, F. 2005. Psycholinguistics, formal grammars, and cognitive science. *The Linguistic Review* 22: 365–80.

- Ferreira, F., Christianson, K., & Hollingworth, A. 2001. Misinterpretations of garden-path sentences: implications for models of reanalysis. *Journal of Psycholinguistic Research* 30: 3–20.
- Fiengo, R. 1977. On trace theory. *Linguistic Inquiry* 8: 35–61.
- Finger, S. 1994. *Origins of neuroscience*. Oxford: Oxford University Press.
- Fisher, S. E. 2006. Tangled webs: tracing the connections between genes and cognition. *Cognition* 101: 270–97.
- Fisher, S. E., Lai, C. S. L., and Monaco, A. P. 2003. Deciphering the genetic basis of speech and language disorders. *Annual Review of Neuroscience* 26: 57–80.
- Fitch, W. T. 2005. The evolution of language: a comparative review. *Biology and Philosophy* 20: 193–230.
- Fitch, W. T. & Hauser, M. D. 2004. Computational constraints on syntactic processing in a nonhuman primate. *Science* 303: 377–80.
- Fitch, W. T., Hauser, M. D., & Chomsky, N. 2005. The evolution of the language faculty: clarifications and implications. *Cognition* 97: 179–210.
- Fodor, J. A. 1975. *The language of thought*. Cambridge, MA: Harvard University Press.
- Fodor, J. A. 1983. *The modularity of mind: an essay in faculty psychology*. Cambridge, MA: MIT Press.
- Fodor, J. A. 1999. Let your brain alone. *London Review of Books*, 14 October.
- Fodor, J. A. 2000. *The mind doesn't work that way: the scope and limits of computational psychology*. Cambridge, MA: MIT Press.
- Fodor, J. A. 2003. *Hume variations*. Oxford: Oxford University Press.
- Fodor, J. A. 2008. *LOT2: the language of thought revisited*. Oxford: Oxford University Press.
- Fodor, J. A. & Pylyshyn, Z. 1988. Connectionism and cognitive architecture: a critical analysis. *Cognition* 28: 3–71.
- Fodor, J. D. 1998. Unambiguous triggers. *Linguistic Inquiry* 29: 1–36.
- Fodor, J. D. 2001. Setting syntactic parameters. In M. Baltin & C. Collins, eds., *The handbook of contemporary syntactic theory* (pp. 730–67). Oxford: Blackwell.
- Fodor, J. D., Bever, T. G., & Garrett, M. F. 1974. *The psychology of language: an introduction to psycholinguistics and generative grammar*. New York: McGraw-Hill.
- Fox, D. & Hackl, M. 2006. The universal density of measurement. *Linguistics and Philosophy* 29: 537–86.
- Fox, M. 2007. *Talking hands*. New York: Simon & Schuster.
- Freidin, R. & Vergnaud, J.-R. 2001. Exquisite connections: some remarks on the evolution of linguistic theory. *Lingua* 111: 639–66.
- Friedlander, B. Z. 1970. Receptive language development in infancy: issues and problems. *Merrill-Palmer Quarterly* 16: 7–51.
- Friedmann, N. 2006. Speech production in Broca's agrammatic aphasia: syntactic tree pruning. In Y. Grodzinsky & K. Amunts, eds., *Broca's region* (pp. 63–82). Oxford: Oxford University Press.
- Frisch, K. von 1967. *The dance language and orientation of bees*. Cambridge, MA: Harvard University Press.
- Frisch, K. von 1974. *Animal architecture*. New York: Harcourt Brace Jovanovich.
- Fromkin, V., ed. 1973. *Speech errors as linguistic evidence*. The Hague/Paris: Mouton.
- Fromkin, V., Krashen, S., Curtiss, S., Rigler, D., & Rigler, M. 1974. The development of language in Genie: a case of language acquisition beyond the "critical period." *Brain and Language* 1: 81–107.

- Fukui, N. 1996. On the nature of economy in language. *Ninti Kagaku [Cognitive Studies]* 3: 51–71.
- Gallistel, C. R. 1990. *The organization of learning*. Cambridge, MA: MIT Press.
- Gallistel, C. R. 2001. Mental representations; psychology of. In N. J. Smelser & P. B. Baltes, eds., *International encyclopedia of the behavioral and social sciences* (pp. 9691–5). Oxford: Elsevier.
- Gallistel, C. R. 2006. The nature of learning and the functional architecture of the brain. In Q. Jin, ed., *Psychological science around the world, vol. 1: proceedings of the 28th International Congress of Psychology* (pp. 63–71). Sussex: Psychology Press.
- Gallistel, C. R. 2007. Learning organs. [English original of L'apprentissage de matières distinctes exige des organes distincts.] In J. Bricmont & J. Franck, eds., *Cahier no 88: Noam Chomsky* (pp. 181–7). Paris: L'Herne.
- Gallistel, C. R. 2008. Learning and representation. In R. Menzel, ed., *Learning and memory: a comprehensive reference* (pp. 227–42). New York: Elsevier.
- Gallistel, C. R. & Gelman, R. 1992. Preverbal and verbal counting and computation. *Cognition* 44: 43–74.
- Gallistel, C. R. & King, A. P. 2009. *Memory and the computational brain: how cognitive science will transform neuroscience*. Malden: Wiley-Blackwell.
- Gambell, T. & Yang, C. D. 2004a. Statistical learning and Universal Grammar: modeling word segmentation. In W. G. Sakas, ed., *Proceedings of the 20th International Conference on Computational Linguistics* (pp. 51–4). Geneva.
- Gambell, T. & Yang, C. D. 2004b. Mechanisms and constraints in word segmentation. Ms., Yale University.
- Gardner, H. 1985. *The mind's new science: a history of the cognitive revolution*. New York: Basic Books.
- Gentner, T. Q., Fenn, K. M., Margoliash, D., & Nusbaum, H. C. 2006. Recursive syntactic pattern learning by songbirds. *Nature* 440: 1204–7.
- Geschwind, N. 1970. The organization of language and the brain. *Science* 170: 940–4.
- Gibson, E. & Wexler, K. 1994. Triggers. *Linguistic Inquiry* 25: 407–54.
- Gleason, J. B. 1958. The child's learning of English morphology. *Word* 14: 150–77.
- Gleitman, L. R. 1990. The structural sources of verb meanings. *Language Acquisition* 1: 3–55.
- Gleitman, L. R. & Fisher, C. 2005. Universal aspects of word learning. In J. McGilvray, ed., *The Cambridge companion to Chomsky* (pp. 123–42). Cambridge: Cambridge University Press.
- Gleitman, L. R. & Landau, B. 1985. *Language and experience: evidence from the blind child*. Cambridge, MA: Harvard University Press.
- Goldin-Meadow, S. 2003. *Hearing gesture: how our hands help us think*. Cambridge, MA: Harvard University Press.
- Goodluck, H. 1991. *Language acquisition: a linguistic introduction*. Oxford: Blackwell.
- Gopnik, M. 1990. Dysphasia in an extended family. *Nature* 344: 715.
- Gordon, R. G., Jr., ed. 2005. *Ethnologue: languages of the world*, 15th edn. Dallas: SIL International.
- Goudarzi, S. 2006. Chomsky challenge: bird said to grasp key aspect of human language. *LiveScience.com*, 26 April.
- Gould, J. L. & Gould, C. G. 1988. *The honey bee*. New York: Freeman.
- Gould, S. J. 2002. *The structure of evolutionary theory*. Cambridge, MA: Belknap.

- Gouteux, S. & Spelke, E. S. 2001. Children's use of geometry and landmarks to reorient in an open space. *Cognition* 81: 119–48.
- Gray, J. 2000. *Hilbert's challenge*. Oxford: Oxford University Press.
- Greenberg, J. H. 1966. Some universals of grammar with particular reference to the meaningful elements. In J. H. Greenberg, ed., *Universals of language* (pp. 73–133). Cambridge, MA: MIT Press.
- Grodzinsky, Y. 1990. *Theoretical perspectives on language development*. Cambridge, MA: MIT Press.
- Grodzinsky, Y. & Amunts, K., eds. 2006. *Broca's region*. Oxford: Oxford University Press.
- Guasti, M. T. 2002. *Language acquisition: the growth of grammar*. Cambridge, MA: Bradford/MIT Press.
- Haegeman, L. 1994. *Introduction to government and binding theory*, 2nd edn. Oxford: Blackwell.
- Haegeman, L. 2006. *Thinking syntactically*. Oxford: Blackwell.
- Haesler, S., Rochefort, C., Georgi, B., Licznarski, P., Osten, P., & Scharff, C. 2007. Incomplete and inaccurate vocal imitation after knockdown of Foxp2 in songbird basal ganglia nucleus area X. *Public Library of Science, Biology* 5: e321.
- Hale, K. & Keyser, S. J. 1993. On argument structure and the lexical expression of syntactic relations. In K. Hale & S. J. Keyser, eds., *The view from building 20* (pp. 53–110). Cambridge, MA: MIT Press.
- Hale, K. & Keyser, S. J. 2002. *Prolegomena to a theory of argument structure*. Cambridge, MA: MIT Press.
- Halle, M. 1978. Knowledge unlearned and untaught: what speakers know about the sounds of their language. In M. Halle, J. Bresnan, & G. A. Miller, eds., *Linguistic theory and psychological reality* (pp. 294–303). Cambridge, MA: MIT Press. (Reprinted in M. Halle, *From speech to language and back: papers on phonetics and phonology, 1954–2002* (pp. 95–104). Berlin: Mouton de Gruyter.)
- Halle, M. 1990. An approach to morphology. In *Proceedings of NELS 20* (pp. 150–84). University of Massachusetts, Amherst: GLSA.
- Halle, M. 1995. Feature geometry and feature spreading. *Linguistic Inquiry* 26: 1–46.
- Halle, M. 2002. *From speech to language and back: papers on phonetics and phonology, 1954–2002*. Berlin: Mouton de Gruyter.
- Halle, M. & Richards, N. 2007. *A biographical memoir of Ken Hale*. Washington, DC: National Academy of Sciences.
- Harrison, P. 1992. Descartes on animals. *The Philosophical Quarterly* 42: 219–27.
- Hauser, M. D. 1996. *The evolution of communication*. Cambridge, MA: MIT Press.
- Hauser, M. D. 2000. *Wild minds: what animals really think*. New York: Henry Holt.
- Hauser, M. D. 2006. *Moral minds*. New York: Ecco.
- Hauser, M. D. 2008. Humaniqueness. Paper presented at the AAAS Annual Meeting, Boston.
- Hauser, M. D. & Carey, S. 2003. Spontaneous representations of small numbers of objects by rhesus macaques: examinations of content and format. *Cognitive Psychology* 47: 367–401.
- Hauser, M. D., Carey, S., & Hauser, L. B. 2000. Spontaneous number representation in semi-free-ranging rhesus monkeys. *Proceedings of the Royal Society: Biological Sciences* 267: 829–33.
- Hauser, M. D., Chomsky, N., & Fitch, W. T. 2002. The faculty of language: what is it, who has it, and how did it evolve? *Science* 298: 1569–79.

- Heim, I. & Kratzer, A. 1998. *Semantics in generative grammar*. Oxford: Blackwell.
- Hermer, L. & Spelke, E. S. 1994. A geometric process for spatial reorientation in young children. *Nature* 370: 57–9.
- Hermer, L. & Spelke, E. S. 1996. Modularity and development: the case of spatial reorientation. *Cognition* 61: 195–232.
- Hermer-Vasquez, L., Moffett, A., & Munkholm, P. 2001. Language, space, and the development of cognitive flexibility in humans. *Cognition* 79: 263–99.
- Hermer-Vasquez, L., Spelke, E. S., & Katsnelson, A. 1999. Sources of flexibility in human cognition: dual-task studies of space and language. *Cognitive Psychology* 39: 3–36.
- Hickock, G. & Poeppel, D. 2000. Towards a functional anatomy of speech processing. *Trends in Cognitive Sciences* 4: 131–8.
- Higginbotham, J. 1995. *Sense and syntax*. Oxford: Oxford University Press.
- Higginbotham, J. & May, R. 1981. Questions, quantifiers, and crossing. *The Linguistic Review* 1: 41–80.
- Hinzen, W. 2006. *Mind design and minimal syntax*. Oxford: Oxford University Press.
- Hinzen, W. 2007. *An essay on names and truth*. Oxford: Oxford University Press.
- Hobbes, T. 1651. *Leviathan*. London: Andrew Crooke.
- Hochmann, J. R., Azadpour, M., & Mehler, J. 2008. Do humans really learn $A^n B^n$ artificial grammars from exemplars? *Cognitive Science* 32: 1021–36.
- Hodges, A. 2008. What did Alan Turing mean by machine? In P. Husbands, O. Holland, & M. Wheeler, eds., *The mechanical mind in history* (pp. 75–90). Cambridge, MA: MIT Press.
- Hoffman, D. D. 1998. *Visual intelligence: how we create what we see*. New York: W. W. Norton.
- Holy, T. E. & Guo, Z. 2005. Ultrasonic songs of male mice. *Public Library of Science, Biology* 3: e386.
- Honda, M. & O’Neil, W. 1993. Triggering science-forming capacity through linguistic inquiry. In K. Hale & S. J. Keyser, eds., *The view from building 20* (pp. 229–55). Cambridge, MA: MIT Press.
- Honda, M. & O’Neil, W. 2008. *Thinking linguistically*. Malden, MA: Blackwell.
- Hornstein, N. 1984. *Logic as grammar: an approach to meaning in natural language*. Cambridge, MA: MIT Press.
- Hornstein, N. 2005. Empiricism and rationalism as research strategies. In J. McGilvray, ed., *The Cambridge companion to Chomsky* (pp. 145–64). Cambridge: Cambridge University Press.
- Hornstein, N. 2009. *A theory of syntax*. Cambridge: Cambridge University Press.
- Hornstein, N., Nunes, J., & Grohmann, K. K. 2006. *Understanding minimalism*. Cambridge: Cambridge University Press.
- Huang, C.-T. J. 1982. Logical relations in Chinese and the theory of grammar. Doctoral dissertation, MIT.
- Hubel, D. H. 1988. *Eye, brain, and vision*. New York: Scientific American Library.
- Hubel, D. H. & Wiesel, T. N. 1962. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *Journal of Physiology of London* 160: 106–54.
- Hubel, D. H. & Wiesel, T. N. 2004. *Brain and visual perception: the story of a 25-year collaboration*. Oxford: Oxford University Press.
- Humboldt, W. von 1827. Über den Dualis. Read at Berlin, Akademie der Wissenschaften, 26 April.

- Humboldt, W. von 1836 [1999]. *On language. On the diversity of human language construction and its influence on the mental development of the human species*. Cambridge: Cambridge University Press.
- Hume, D. 1739–40. *A treatise of human nature*. London: John Noon & Longman.
- Hume, D. 1748. *An enquiry concerning human understanding*. London: Millar.
- Husbands, P., Holland, O., & Wheeler, M., eds. 2008. *The mechanical mind in history*. Cambridge, MA: MIT Press.
- Jackendoff, R. 1977. *X-bar syntax*. Cambridge, MA: MIT Press.
- Jackendoff, R. 1983. *Semantics and cognition*. Cambridge, MA: MIT Press.
- Jackendoff, R. 1987. *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
- Jackendoff, R. 1989. A comparison of rhythmic structures in music and language. In P. Kiparsky & G. Youmans, eds., *Phonetics and phonology* (vol. 1, pp. 15–44). New York: Academic Press.
- Jackendoff, R. 1990. *Semantic structures*. Cambridge, MA: MIT Press.
- Jackendoff, R. 1992. *Languages of the mind*. Cambridge, MA: Bradford/MIT Press.
- Jackendoff, R. 1994. *Patterns in the mind: language and human nature*. New York: Basic Books.
- Jackendoff, R. 1997. *The architecture of the language faculty*. Cambridge, MA: MIT Press.
- Jackendoff, R. 2002. *Foundations of language: brain, meaning, grammar, evolution*. Oxford: Oxford University Press.
- Jackendoff, R. 2007. *Language, consciousness, culture: essays on mental structure*. Cambridge, MA: MIT Press.
- Jackendoff, R. & Lerdahl, F. 2006. The capacity for music: what's special about it? *Cognition* 100: 33–72.
- Jackendoff, R. & Pinker, S. 2005. The nature of the language faculty and its implications for evolution of language (Reply to Fitch, Hauser, & Chomsky). *Cognition* 97: 211–25.
- Jacob, F. 1977. Evolution and tinkering. *Science* 196: 1161–6.
- Jacob, F. 1982. *The possible and the actual*. New York: Pantheon.
- Jacob, F. & Monod, J. 1961. Genetic regulatory mechanisms in the synthesis of proteins. *Journal of Molecular Biology* 3: 318–56.
- James, W. 1887. What is an instinct? *Scribner's Magazine* 1: 355–65.
- James, W. 1890. *The principles of psychology*. New York: Holt.
- Jenkins, L. 2000. *Biolinguistics: exploring the biology of language*. Cambridge: Cambridge University Press.
- Jerne, N. K. 1985. The generative grammar of the immune system. *Science* 229: 1057–9.
- Jones, P., ed. 1989. *The science of man in the Scottish Enlightenment: Hume, Reid, and their contemporaries*. Edinburgh: John Donald.
- Joos, M., ed. 1957. *Readings in linguistics*. New York: American Council of Learned Societies.
- Jusczyk, P. 1997. *The discovery of spoken language*. Cambridge, MA: MIT Press/Bradford Books.
- Kandel, E. R., Schwartz, J., & Jessell, T. 2000. *Principles of neural science*, 4th edn. New York: McGraw-Hill.
- Kaufman, E. L., Lord, M. W., Reese, T. W., & Volkman, J. 1949. The discrimination of visual number. *American Journal of Psychology* 62: 498–525.
- Kayne, R. 1975. *French syntax*. Cambridge, MA: MIT Press.
- Kayne, R. 1984. *Connectedness and binary branching*. Dordrecht: Foris.
- Kayne, R. 1994. *The antisymmetry of syntax*. Cambridge, MA: MIT Press.

- Kayne, R. 2000. *Parameters and universals*. Oxford: Oxford University Press.
- Kayne, R. 2005. *Movement and silence*. Oxford: Oxford University Press.
- Keats, J. 1819. *Lamia*. (Poem included in 1884, *The Poetical Works of John Keats*.)
- Keenan, E. & Stavi, J. 1986. A semantic characterization of natural language determiners. *Linguistics and Philosophy* 9: 253–326.
- Kegl, J. 2004. Language emergence in a language-ready brain: acquisition issues. In G. Morgan & B. Woll, eds., *Language acquisition in signed languages* (pp. 207–54). Cambridge: Cambridge University Press.
- Kegl, J., Senghas, A., & Coppola, M. 1999. Creation through contact: sign language emergence and sign language change in Nicaragua. In M. DeGraff, ed., *Comparative grammatical change: the intersection of language acquisition, creole genesis, & diachronic syntax* (pp. 179–237). Cambridge, MA: MIT Press.
- Keyser, S. J. 1970. The role of linguistics in the elementary school curriculum. *Elementary English* 47: 39–45.
- Keyser, S. J. 2001. Appreciation. In M. Kenstowicz, ed., *Ken Hale: a life in language* (pp. xi–xiii). Cambridge, MA: MIT Press.
- Kingsley, D. 1940. Extreme isolation of a child. *American Journal of Sociology*, January issue, 554–65.
- Kingsley, D. 1947. Final note on a case of extreme isolation. *American Journal of Sociology*, March issue, 432–7.
- Kinzler, K. D. & Spelke, E. S. 2005. The effect of language on infants' preference for faces. Poster presented at the Biennial Meeting of the Society for Research in Child Development, Atlanta.
- Kinzler, K. D. & Spelke, E. S. 2007. Core systems in human cognition. *Progress in Brain Research* 164: 257–64.
- Koster, J. 1978. *Locality principles in syntax*. Dordrecht: Foris.
- Koster, J., Riemsdijk, H. van, & Vergnaud, J.-R. 1978. The GLOW Manifesto. In *The First GLOW Newsletter*.
- Kroch, A. 2000. Syntactic change. In M. Baltin & C. Collins, eds., *The handbook of contemporary syntactic theory* (pp. 629–739). Malden, MA: Blackwell.
- Ladusaw, W. 1979. Polarity sensitivity as inherent scope relations. Doctoral dissertation, University of Texas – Austin.
- Lai, C. S. L., Fisher, S. E., Hurst, J. A., Vargha-Khadem, F., & Monaco, A. P. 2001. A forkhead-domain gene is mutated in a severe speech and language disorder. *Nature* 413: 519–23.
- Lane, H. 1976. *The wild boy of Aveyron*. Cambridge, MA: Harvard University Press.
- Larson, R. 1996. Grammar as a laboratory science. Paper presented at the AAAS Annual Meeting, Baltimore.
- Larson, R. & Segal, G. 1995. *Knowledge of meaning*. Cambridge, MA: MIT Press.
- Lashley, K. S. 1951. The problem of serial order in behavior. In L.A. Jeffress, ed., *Cerebral mechanisms in behavior* (pp. 112–36). New York: Wiley.
- Lasky, R. E., Syrdal-Lasky, A., & Klein, R. E. 1975. VOT discrimination by four to six and a half month old infants from Spanish environments. *Journal of Experimental Child Psychology* 20: 213–25.
- Lasnik, H. 2000. *Syntactic structures revisited*. Cambridge, MA: MIT Press.
- Lasnik, H. 2005. Grammar, levels, and biology. In J. McGilvray, ed., *The Cambridge companion to Chomsky* (pp. 60–83). Cambridge: Cambridge University Press.

- Lasnik, H. & Saito, M. 1984. On the nature of proper government. *Linguistic Inquiry* 15: 235–89.
- Lasnik, H., Uriagereka, J., & Boeckx, C. 2005. *A course in minimalist syntax*. Oxford: Blackwell.
- Legate, J. A. & Yang, C. D. 2002. Empirical re-assessment of stimulus poverty arguments. *The Linguistic Review* 19: 151–62.
- Lehrer, J. 2007. *Proust was a neuroscientist*. New York: Houghton Mifflin.
- Leibniz, G. 1686. *Discours de métaphysique*. Trans. In *Leibniz: discourse on metaphysics and correspondence with Arnauld*, ed. G. R. Montgomery, 1902. La Salle, IL.: Open Court.
- Leibniz, G. 1982. *New essays on human understanding*, ed. P. Remnant & J. Bennett. Cambridge: Cambridge University Press.
- Lely, H. K. J. van der 2005. Domain-specific cognitive systems: insight from grammatical specific language impairment. *Trends in Cognitive Science* 9: 53–9.
- Lenneberg, E. H. 1960. Language, evolution, and purposive behavior. In S. Diamond, ed., *Culture in history* (pp. 869–93). New York: Columbia University Press. (Extended version published as “On the capacity of language acquisition”, 1964, in J. Fodor and J. Katz, eds., *The structure of language* (pp. 579–603). Englewood Cliffs, NJ: Prentice-Hall.)
- Lenneberg, E. H. 1964. A biological perspective of language. In E. H. Lenneberg, ed., *New directions in the study of language* (pp. 65–88). Cambridge, MA: MIT Press.
- Lenneberg, E. H. 1966. The natural history of language. In F. Smith & G. A. Miller, eds., *The genesis of language: a psycholinguistic approach* (pp. 219–52). Cambridge, MA: MIT Press.
- Lenneberg, E. H. 1967. *Biological foundations of language*. New York: Wiley.
- Lerdahl, F. & Jackendoff, R. 1983. *A generative theory of tonal music*. Cambridge, MA: MIT Press.
- Lewontin, R. C. 1998. The evolution of cognition: questions we will never answer. In D. Scarborough, S. Sternberg, & D. N. Osherson, eds., *An invitation to cognitive science* (pp. 107–32). Cambridge, MA: MIT Press.
- Lewontin, R. C. 2000a. *The triple helix: gene, organism, and environment*. Cambridge, MA: Harvard University Press.
- Lewontin, R. C. 2000b. *It ain't necessarily so: the dream of the human genome and other illusions*. Cambridge, MA: Harvard University Press.
- Liberman, M. 2006. Starlings. <http://itre.cis.upenn.edu/~myl/languagelog/archives/003076.html>.
- Lichtheim, L. 1885. On aphasia. *Brain* 7: 433–84.
- Lidz, J. & Pearl, L. 2007. Language learning through selection, review of C. Yang, *The infinite gift: how children learn and unlearn the languages of the world*. *Science* 315: 332.
- Lightfoot, D. 1982. *The language lottery: toward a biology of grammars*. Cambridge, MA: MIT Press.
- Lightfoot, D. 1989. The child's trigger experience: degree-0 learnability. *Behavioral and Brain Sciences* 12: 321–75.
- Lightfoot, D. 1991. *How to set parameters: arguments from language change*. Cambridge, MA: MIT Press.
- Lightfoot, D. 1999. *The development of language: acquisition, change, and evolution*. Oxford: Blackwell.
- Lightfoot, D. 2006. *How new languages emerge*. Cambridge: Cambridge University Press.

- Lipton, J. S. & Spelke, E. S. 2003. Origins of number sense: large number discrimination in human infants. *Psychological Science* 14: 396–401.
- Locke, J. 1690. *An essay concerning human understanding*. London.
- Longobardi, G. 2003. Methods in parametric linguistics and cognitive history. *Linguistic Variation Yearbook* 3: 101–38.
- Lorenz, K. 1937. Über die Bildung des Instinktsbegriffes. *Die Naturwissenschaften* 25: 289–300, 307–18, 325–31.
- Lorenz, K. 1939. Vergleichende Verhaltensforschung. *Zoologischer Anzeiger*, Suppl. 12: 69–102.
- Lorenz, K. 1970. *Studies in animal and human behavior*, vol. 1. Cambridge, MA: Harvard University Press.
- Lorenz, K. 1971. *Studies in animal and human behavior*, vol. 2. Cambridge, MA: Harvard University Press.
- Lorenzo, G. & Longa, V. M. 2003. *Homo loquens*. Lugo: Tris Tram.
- Ludlow, P. 2002. LF and natural logic. In G. Preyer & G. Peter, eds., *Logical form and language* (pp. 132–68). Oxford: Oxford University Press.
- Luria, S. 1973. *Life: the unfinished experiment*. New York: Scribner.
- Macnamara, J. 1999. *Through the rearview mirror: historical reflections on psychology*. Cambridge, MA: MIT Press.
- Malebranche, N. 1958–78. *Oeuvres complètes* (20 vols.). Paris: J. Vrin.
- Malson, L. 1972. *Wolf children and the problem of human nature*. New York: Monthly Review Press.
- Mandelbrot, B. 1982. *The fractal geometry of nature*. San Francisco: Freeman.
- Marantz, A. 2005. Generative linguistics within the cognitive neuroscience of language. *The Linguistic Review* 22: 429–45.
- Marcus, G. F. 2001. *The algebraic mind: integrating connectionism and cognitive science*. Cambridge, MA: MIT Press.
- Marcus, G. F. 2004. *The birth of the mind: how a tiny number of genes creates the complexities of human thought*. New York: Basic Books.
- Marcus, G. F. 2006. Cognitive architecture and descent with modification. *Cognition* 101: 443–65.
- Marcus, G. F. & Fisher, S. E. 2003. FOXP2 in focus: what can genes tell us about speech and language? *Trends in Cognitive Sciences* 7: 257–62.
- Marcus, G. F., Pinker, S., Ullman, M., Hollander, M., Rosen, T. J., & Xu, F. 1992. *Over-regularization in language acquisition* (Monographs of the Society for Research in Child Development 57). Chicago: University of Chicago Press.
- Margules, J. & Gallistel, C. R. 1988. Heading in the rat: determination by environmental shape. *Animal Learning and Behavior* 16: 404–10.
- Marr, D. 1982. *Vision: a computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.
- Mayr, E. 1970. *Populations, species, and evolution: an abridgement of animal species and evolution*. Cambridge, MA: Harvard University Press.
- Mayr, E. 1982. *The growth of biological thought*. Cambridge: Belknap.
- McCloskey, J. 2004. Unravelling the threads. Paper presented at the Irish Association for Applied Linguistics Annual Meeting, Dublin.
- McDaniel, D. 1986. Conditions on *wh*-chains. Doctoral dissertation, CUNY.
- McGilvray, J. 1999. *Chomsky: language, mind, and politics*. Cambridge: Polity.

- McGivray, J., ed. 2005. *The Cambridge companion to Chomsky*. Cambridge: Cambridge University Press.
- McGinn, C. 1989. Can we solve the mind–body problem? *Mind* 97: 349–66.
- McNeil, D. 1966. Developmental psycholinguistics. In G. A. Miller, ed., *The genesis of language* (pp. 15–84). Cambridge, MA: MIT Press.
- Mehler, J. 1974. Apprendre par désapprendre. In *L'unité de l'homme, le cerveau humain*. Paris.
- Mehler, J., and Dupoux, E. 1994. *What infants know*. Oxford: Blackwell.
- Meroni, L., Gualmini, A., & Crain, S. 2001. Three years of continuous acquisition. *Proceedings of the Second Tokyo Conference on Psycholinguistics* (pp. 1–33). Tokyo: Hituzi.
- Mikhail, J. 2000. Rawls' linguistic analogy: a study of the “generative grammar” model of moral theory described by John Rawls in *A theory of justice*. Doctoral dissertation, Cornell.
- Mikhail, J. 2007. Universal moral grammar: theory, evidence, and the future. *Trends in Cognitive Sciences* 11: 143–52.
- Mikhail, J. To appear. *Moral grammar: Rawls' linguistic analogy and the cognitive science of moral judgment*. Cambridge: Cambridge University Press.
- Mill, J. S. 1843. *A system of logic*. London: Longman.
- Miller, G. A. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review* 63: 81–97.
- Miller, G. A. 1998. Interview. In M. Gazzaniga, R. Ivry, and G. Mangun, eds., *Fundamentals of cognitive neuroscience*. New York: Norton.
- Miller, G. A. 2003. The cognitive revolution: a historical perspective. *Trends in Cognitive Sciences* 7: 141–4.
- Mithen, S. 1996. *The prehistory of the mind: the cognitive origins of art, religion, and science*. London: Thames & Hudson.
- Montague, R. 1970a. Universal grammar. *Theoria* 36 (1970): 373–98. (Reprinted 1974 in R. Thomason, ed., *Formal philosophy: selected papers by Richard Montague*. New Haven: Yale University Press.)
- Montague, R. 1970b. English as a formal language. In B. Visentini (ed.), *Linguaggi nella società e nella tecnica* (pp. 189–223). Milan: Edizioni di Comunità. (Reprinted 1974 in R. Thomason, ed., *Formal philosophy: selected papers by Richard Montague*. New Haven: Yale University Press.)
- Montague, R. 1973. The proper treatment of quantification in ordinary English. In K. J. J. Hintikka, J. M. E. Moravcsik, & P. Suppes, eds., *Approaches to natural language* (pp. 221–42). Dordrecht: Reidel.
- Moon, C., Cooper, R. P., & Fifer, W. P. 1993. Two-day-olds prefer their native language. *Infant Behavior and Development* 16: 495–500.
- Moravcsik, J. 1998. *Meaning, creativity, and the partial inscrutability of the human mind*. Stanford: CSLI.
- Morgan, J. & Demuth, K., eds. 1996. *Signal to syntax: bootstrapping from speech to grammar in early acquisition*. Mahwah, NJ: Lawrence Erlbaum.
- Moro, A. 2000. *Dynamic antisymmetry*. Cambridge, MA: MIT Press.
- Moro, A. 2008. *The boundaries of Babel: the brain and the enigma of impossible languages*. Cambridge, MA: MIT Press.
- Moro, A., Tettamanti, M., Perani, D., Donati, C., Cappa, S. C., & Fazio, F. 2001. Syntax and the brain: disentangling grammar by selective anomalies. *NeuroImage* 13: 110–18.

- Musso, M., Moro, A., Glauche, V., Rijntjes, M., Reichenbach, J., Büchel, C., & Weiller, C. 2003. Broca's area and the language instinct. *Nature Neuroscience* 6: 774–81.
- Nagel, T. 1993. The mind wins! *New York Review of Books*, March 4. (Reprinted as "Searle: why we are not computers" in T. Nagel, *Other minds*. Oxford: Oxford University Press.)
- Napoli, D. J. 2003. *Language matters: a guide to everyday questions about language*. Oxford: Oxford University Press.
- Neisser, U. 1967. *Cognitive psychology*. New York: Appleton-Century-Crofts.
- Nespor, M. 2001. About parameters, prominence, and bootstrapping. In E. Dupoux, ed., *Language, brain, and cognitive development: essays in honor of Jacques Mehler* (pp. 127–42). Cambridge, MA: MIT Press.
- Nevins, A. 2004. What UG can and can't do to help the reduplication learner. In *Plato's problems: papers on language acquisition – MITWPL* 48: 113–26.
- Newmeyer, F. 1998. *Language form and language function*. Cambridge, MA: MIT Press.
- Newmeyer, F. 2005. *Possible and probable languages*. Oxford: Oxford University Press.
- Newton, M. 2002. *Savage boys and wild girls: a history of feral children*. London: Faber & Faber.
- Niyogi, P. 2006. *The computational nature of language learning and evolution*. Cambridge, MA: MIT Press.
- Ochs, E. 1982. Talking to children in Western Samoa. *Language and Society* 2: 77–104.
- O'Grady, W. 2005. *How children learn language*. Cambridge: Cambridge University Press.
- O'Kane, G. C. & Spelke, E. S. 2001. Language dependence in representations of number, space, and time. Poster presented at the meeting of the Psychonomic Society, Orlando.
- O'Neil, W. 1969. Foreword to N. R. Cattell, *The new English grammar: a descriptive introduction*. Cambridge, MA: MIT Press.
- Otero, C. P. 1990. The emergence of *homo loquens* and the laws of physics. *Behavioral and Brain Sciences* 13: 747–50.
- Papanicolaou, A. C. 1998. *Fundamentals of functional brain imaging*. Lisse: Swets & Zeitlinger.
- Partee, B. H. 1979. Semantics: mathematics or psychology? In R. Bauerle, U. Egli, & A. von Stechow, eds., *Semantics from different points of view* (pp. 1–14). Berlin: Springer-Verlag.
- Patel, A. 2008. *Language and music*. Oxford: Oxford University Press.
- Pavlov, I. P. 1927. *Conditioned reflexes: an investigation of the physiological activity of the cerebral cortex*. London: Oxford University Press.
- Pearl, L. 2007. Necessary bias in natural language learning. Doctoral dissertation, University of Maryland – College Park.
- Peirce, C. S. 1931–66. *The collected papers of C.S. Peirce*. Cambridge, MA: Harvard University Press.
- Penfield, W. & Rasmussen, T. 1950. *The cerebral cortex of man*. New York: Macmillan.
- Pernau, J. F. A. von 1716. *Unterricht, was mit dem lieblichen Geschoepff denen Voegeln, auch ausser dem Fang, nur durch due Ergruendung deren Eigenschaften, und Zahmmachung, oder anderer Abrichtung, Man sich von Lust und Zeit-Vertrieb machen koennen*. Nurenberg.
- Pettit, P. 2008. *Hobbes on language, mind, and politics*. Princeton, NJ: Princeton University Press.
- Petitto, L. A. 2005. How the brain begets language: on the neural tissue underlying human language acquisition. In J. McGilvray, ed., *The Cambridge companion to Chomsky* (pp. 84–101). Cambridge: Cambridge University Press.

- Petitto, L. A., Holowka, S., Sergio, L., & Ostry, D. 2001. Language rhythms in babies' hand movements. *Nature* 413: 35–6.
- Petzold, C. 2008. *The annotated Turing*. New York: Wiley.
- Phillips, C. 1996. Order and structure. Doctoral dissertation, MIT.
- Phillips, C. 2004. Linguistics and linking problems. In M. Rice & S. Warren, eds., *Developmental language disorders: from phenotypes to etiologies* (pp. 241–87). Mahwah, NJ: Lawrence Erlbaum.
- Phillips, C. 2006. Three benchmarks for distributional models of syntax. In R. Zanuttini, H. Campos, E. Herburger, & P. Portner, eds., *Negation, tense, and clausal architecture: cross-linguistic investigations* (pp. 13–26). Washington, DC: Georgetown University Press.
- Phillips, C. To appear. Should we impeach armchair linguists? In S. Iwasaki, ed., *Japanese-Korean Linguistics* 17. Stanford: CSLI.
- Phillips, C. & Lau, E. 2004. Foundational issues (review of R. Jackendoff, *Foundations of language*). *Journal of Linguistics* 40: 571–91.
- Phillips, C. & Sakai, K. L. 2005. Language and the brain. In *The McGraw-Hill handbook of science and technology* (pp. 166–9). New York: McGraw-Hill.
- Phillips, C. & Wagers, M. 2007. Relating structure and time in linguistics and psycholinguistics. In G. Gaskell, ed., *The Oxford handbook of psycholinguistics* (pp. 739–56). Oxford: Oxford University Press.
- Piaget, J. 1928. *The child's conception of the world*. London: Routledge & Kegan Paul.
- Piaget, J. 1952. *The child's conception of number*. London: Routledge & Kegan Paul.
- Piaget, J. 1955. *The child's construction of reality*. London: Routledge & Kegan Paul.
- Piattelli-Palmarini, M., ed. 1980. *Language and learning: the debate between Jean Piaget and Noam Chomsky*. Cambridge, MA: Harvard University Press.
- Piattelli-Palmarini, M. 1987. The rise of selective theories: a case study and some lessons from immunology. In A. Marras & W. Demopoulos, eds., *Language learning and concept acquisition: foundational issues* (pp. 117–30). Norwood, NJ: Ablex.
- Piattelli-Palmarini, M. 1989. Evolution, selection, and cognition: from “learning” to parameter setting in biology and in the study of language. *Cognition* 31: 1–44.
- Piattelli-Palmarini, M. 1994. Ever since language and learning: afterthoughts on the Piaget–Chomsky debate. *Cognition* 50: 315–46.
- Piattelli-Palmarini, M. 2008. *Le scienze cognitive classiche: un panorama*. Torino: Einaudi Editore.
- Piattelli-Palmarini, M. & Uriagereka, J. 2005. The evolution of the narrow faculty of language: the skeptical view and a reasonable conjecture. *Lingue e Linguaggio* 4: 27–79.
- Piattelli-Palmarini, M. & Uriagereka, J. To appear. A geneticist's dream, a linguist's nightmare: the case of FOXP2. In A.-M. Di Sciullo & C. Boeckx, eds., *The biolinguistic enterprise*. Oxford: Oxford University Press.
- Pickover, C. A. 2006. *The Möbius strip: Dr. August Möbius's marvelous band in mathematics, games, literature, art, technology, and cosmology*. New York: Thunder's Mouth Press.
- Pietroski, P. 2003. The character of natural language semantics. In A. Barber, ed., *Epistemology of language* (pp. 217–56). Oxford: Oxford University Press.
- Pietroski, P. 2005a. Meaning before truth. In G. Preyer and G. Peters, eds., *Contextualism in philosophy* (pp. 253–300). Oxford: Oxford University Press.

- Pietroski, P. 2005b. *Events and semantic architecture*. Oxford: Oxford University Press.
- Pietroski, P. 2006. Character before content. In J. Thomson & A. Byrne, eds., *Content and modality: themes from the philosophy of Robert Stalnaker* (pp. 34–60). Oxford: Oxford University Press.
- Pietroski, P. 2007. Systematicity via monadicity. *Croatian Journal of Philosophy* 7: 343–74.
- Pietroski, P. To appear. *Semantics without truth values*. Oxford: Oxford University Press.
- Pietroski, P., Halberda, J., Lidz, J., & Hunter, T. 2007. Beyond truth conditions: an investigation into the semantics of “most.” Ms., University of Maryland & Johns Hopkins University. [Revised version to appear, *Mind and Language*.]
- Pinker, S. 1984. *Language learnability and language development*. Cambridge, MA: Harvard University Press.
- Pinker, S. 1994. *The language instinct*. New York: Harper Collins.
- Pinker, S. 1997. *How the mind works*. New York: W. W. Norton.
- Pinker, S. 1999. *Words and rules*. New York: Harper Perennial.
- Pinker, S. 2002. *The blank slate: the modern denial of human nature*. New York: Viking.
- Pinker, S. 2008. The moral instinct. *New York Times*, 13 January.
- Pinker, S. & Jackendoff, R. 2005. What’s special about the human language faculty? *Cognition* 95: 201–36.
- Pizlo, Z. 2008. *3D shape: its unique place in visual perception*. Cambridge, MA: MIT Press.
- Plato 1956. *Protagoras and Meno*. New York: Penguin Classics.
- Poeppel, D. 2005. The interdisciplinary study of language and its challenges. In D. Grimm, ed., *Jahrbuch des Wissenschaftskollegs zu Berlin*.
- Poeppel, D. & Embick, D. 2005. The relation between linguistics and neuroscience. In A. Cutler, ed., *Twenty-first century psycholinguistics: four cornerstones* (pp. 173–89). Hillsdale, NJ: Lawrence Erlbaum.
- Poeppel, D. & Hickok, G. 2004. Introduction: towards a new functional anatomy of language. *Cognition* 92: 1–12.
- Pollock, J.-Y. 1989. Verb movement, universal grammar, and the structure of IP. *Linguistic Inquiry* 20: 365–424.
- Portner, P. 2005. *What is meaning? Fundamentals of formal semantics*. Oxford: Blackwell.
- Priestley, J. 1777. *Disquisitions relating to matter and spirit*. London: Printed for J. Johnson.
- Pullum, G. 2004. Ideology, power, and linguistic theory. Ms., University of California, Santa Cruz.
- Pullum, G. & Scholz, B. 2002. Empirical assessment of stimulus poverty arguments. *The Linguistic Review* 19: 9–50.
- Pulvermüller, F. 2003. *The neuroscience of language: on brain circuits of words and serial order*. Cambridge: Cambridge University Press.
- Pustejovsky, J. 1995. *The generative lexicon*. Cambridge, MA: Bradford/MIT Press.
- Putnam, H. 1975. The meaning of “meaning.” In K. Gunderson, ed., *Language, mind, and knowledge* (pp. 131–93). Minneapolis: University of Minnesota Press.
- Pylyshyn, Z. 1984. *Computation and cognition: towards a foundation for cognitive science*. Cambridge, MA: MIT Press.
- Pylyshyn, Z. 1999. Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences* 22: 341–423.
- Quine, W. V. O. 1960. *Word and object*. Cambridge, MA: MIT Press.
- Ramus, F. 1999. Languages’ rhythm and language acquisition. Doctoral dissertation, École des Hautes Études en Sciences Sociales.

- Ramus, F. & Mehler, J. 1999. In praise of functional psychology: dismissing the skeptics. In S. M. Kosslyn, A. M. Galaburda, & Y. Christen, eds., *Languages of the brain* (pp. 897–908). Cambridge, MA: MIT Press.
- Randall, L. 2005. *Warped passages: unraveling the universe's hidden dimensions*. New York: Ecco.
- Rawls, J. 1971. *A theory of justice*. Cambridge, MA: Harvard University Press.
- Reimarus, H. S. 1762. *Allgemeine Betrachtungen über die Triebe der Thiere, hauptsächlich über ihre Kunsttriebe*. Hamburg.
- Rigden, J. S. 2002. *Hydrogen: the essential element*. Cambridge, MA: Harvard University Press.
- Rizzi, L. 1978. Violations of the wh-island constraint in Italian and the subjacency condition. *Montreal Working Papers in Linguistics* 11: 155–90.
- Rizzi, L. 1986. Null objects in Italian and the theory of *pro*. *Linguistic Inquiry* 17: 501–57.
- Rizzi, L. 1989. The new comparative syntax: principles and parameters of Universal Grammar. *Cahiers Ferdinand de Saussure* 43: 65–78.
- Rizzi, L. 1990. *Relativized minimality*. Cambridge, MA: MIT Press.
- Rizzi, L. 1997. The fine structure of the left periphery. In L. Haegeman, ed., *Elements of grammar* (pp. 281–337). Dordrecht: Kluwer.
- Rizzi, L. 2001. Relativized minimality effects. In M. Baltin & C. Collins, eds., *The handbook of contemporary syntactic theory* (pp. 89–110). Oxford: Blackwell.
- Rizzi, L. 2004. On the study of the language faculty: results, developments, and perspectives. *The Linguistic Review* 21: 323–44.
- Roberts, I., ed. 2006. *Comparative grammar*. London: Routledge.
- Ross, J. R. 1967. Constraints on variables in syntax. Doctoral dissertation, MIT. (Published 1986 as *Infinite syntax!* Norwood, NJ: Ablex.)
- Rutter, M. 2006. *Genes and behavior: nature–nurture interplay, explained*. Oxford: Blackwell.
- Rymer, R. 1994. *Genie: a scientific tragedy*. London: HarperPerennial.
- Saffran, J., Aslin, R., & Newport, E. 1996. Statistical learning by 8-month old infants. *Science* 274: 1926–8.
- Saffran, J., Senghas, A., & Trueswell, J. C. 2001. The acquisition of language by children. *Proceedings of the National Academy of Sciences* 98: 12874–5.
- Sakas, W. G. & Fodor, J. D. 2001. The structural triggers learner. In S. Bertolo, ed., *Language acquisition and learnability* (pp. 172–233). Cambridge: Cambridge University Press.
- Sandler, W. & Lillo-Martin, D. 2006. *Sign language and linguistic universals*. Cambridge: Cambridge University Press.
- Sandler, W., Meir, I., Padden, C., & Aronoff, M. 2005. The emergence of grammar in a new sign language. *Proceedings of the National Academy of Sciences* 102: 2661–5.
- Santayana, G. 1905. *The life of reason, vol. 1: reason in common sense*. New York: Scribner.
- Santos, L. R., Hauser, M. D., & Spelke, E. S. 2002. Domain-specific knowledge in human children and non-human primates: artifact and food kinds. In M. Bekoff, ed., *The cognitive animal* (pp. 205–15). Cambridge, MA: MIT Press.
- Sapir, E. 1921. *Language: an introduction to the study of speech*. New York: Harcourt Brace.
- Saussure, F. de 1916. *Cours de linguistique générale*. Paris: Payot.
- Schreyer, R. 1985. The origin of language: a scientific approach to the study of man. *Topoi* 4: 181–6.
- Schreyer, R. 1989. Language and the science of man in the Scottish Enlightenment. In P. Jones, ed., *The science of man in the Scottish Enlightenment* (pp. 149–77). Edinburgh: Edinburgh University Press.

- Schütze, C. 1996. *The empirical base of linguistics: grammaticality judgments and linguistic methodology*. Chicago: University of Chicago Press.
- Senghas, A. 2005. Language emergence: clues from a new Bedouin sign language. *Current Biology* 15: 463–5.
- Shannon, C. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27: 379–423, 623–56.
- Shannon, C. & Weaver, W. 1949. *The mathematical theory of communication*. Urbana, IL: University of Illinois Press.
- Shieber, S. M., ed. 2004. *The Turing test*. Cambridge, MA: MIT Press.
- Shu, W., Cho, J. Y., Jiang, Y. et al. 2005. Altered ultrasonic vocalization in mice with a disruption in the Foxp2 gene. *Proceedings of the National Academy of Science USA* 102: 9643–8.
- Skinner, B. F. 1938. *The behavior of organisms: an experimental analysis*. New York: Appleton-Century-Crofts.
- Skinner, B. F. 1953. *Science and human behavior*. New York: Macmillan.
- Skinner, B. F. 1957. *Verbal behavior*. New York: Appleton-Century-Crofts.
- Smith, N. 2005. Chomsky's science of language. In J. McGilvray, ed., *The Cambridge companion to Chomsky* (pp. 21–41). Cambridge: Cambridge University Press.
- Smith, N. & Tsimpli, I.-M. 1995. *Language learning and modularity: a case-study of a polyglot savant*. Oxford: Blackwell.
- Spelke, E. S. 1988a. Where perceiving ends and thinking begins: the apprehension of objects in infancy. In A. Yonas, ed., *Perceptual development in infancy* (Minnesota Symposium on Child Psychology, vol. 20, pp. 197–234). Hillsdale, NJ: Lawrence Erlbaum.
- Spelke, E. S. 1988b. The origins of physical knowledge. In L. Weiskrantz, ed., *Thought without language*. Oxford: Oxford University Press.
- Spelke, E. S. 1990. Principles of object perception. *Cognitive Science* 14: 29–56.
- Spelke, E. S. 1994. Initial knowledge: Six suggestions. *Cognition*, 50, 431–45.
- Spelke, E. S. 2000. Core knowledge. *American Psychologist* 55: 1233–43.
- Spelke, E. S. 2003a. What makes us smart? Core knowledge and natural language. In D. Gentner and S. Goldin-Meadow, eds., *Language in mind: advances in the investigation of language and thought* (pp. 277–311). Cambridge, MA: MIT Press.
- Spelke, E. S. 2003b. Developing knowledge of space: core systems and new combinations. In S. M. Kosslyn, A. Galaburda, & Y. Christen, eds., *The Languages of the brain* (pp. 239–58). Cambridge, MA: Harvard University Press.
- Spelke, E. S. 2004. Core knowledge. In N. Kanwisher & J. Duncan, eds., *Attention and performance, vol. 20: Functional neuroimaging of visual cognition* (pp. 29–56). Oxford: Oxford University Press.
- Spelke, E. S. & Kinzler, K. D. 2007. Core knowledge. *Developmental Science* 10: 89–96.
- Spelke, E. S. & Tsivkin, S. 2001. Language and number: a bilingual training study. *Cognition* 78: 45–88.
- Spelke, E. S., Breinlinger, K., Macomber, J., & Jacobson, K. 1992. Origins of knowledge. *Psychological Review* 99: 605–32.
- Sperry, R. W. 1968. Plasticity of neural maturation. *Developmental Psychology Supplement* 2: 306–27.
- Spinoza, B. 1677. *Ethica ordine geometrico demonstrata*. In *Opera postuma*.
- Sprouse, J. 2007. A program for experimental syntax. Doctoral dissertation, University of Maryland – College Park.

- Stainton, R. 2006. Meaning and reference: some Chomskyan themes. In E. Lepore and B. Smith, eds., *Oxford handbook of philosophy of language* (pp. 913–40). Oxford: Oxford University Press.
- Strawson, P. F. 1950. On referring. *Mind* 59: 320–44.
- Streeter, L. A. 1976. Language perception of 2-month-old infants shows effects of both innate mechanisms and experience. *Nature*, 259: 39–41.
- Tattersall, I. 1998. *Becoming human: evolution and human uniqueness*. New York: Harcourt Brace.
- Tesan, G. & Thornton, R. 2003. Small children's big grammars. In Y. Otsu, ed., *Proceedings of the Fourth Tokyo Conference on Psycholinguistics* (pp. 239–64). Tokyo: Hitzu.
- Tesan, G. & Thornton, R. 2005. Emergence of do-support in child English. In S. Arunachalam, T. Scheffler, S. Sundaresan, & J. Tauberer, eds., *UPenn Working Papers in Linguistics* 11: 265–78. Philadelphia: Penn Linguistics Club.
- Thornton, R. 1990. Adventures in long-distance moving: the acquisition of complex wh-questions. Doctoral dissertation, University of Connecticut.
- Thornton, R. 1995. Referentiality and wh-movement in child English: juvenile d-linkuency. *Language Acquisition* 4: 139–75.
- Thornton, R. 2008. Why continuity. *Natural Language and Linguistic Theory* 26: 107–46.
- Thornton, R. & Tesan, G. 2007. Categorical acquisition: parameter setting in Universal Grammar. *Biolinguistics* 1: 49–98.
- Tinbergen, N. 1951. *The study of instinct*. Oxford: Clarendon Press.
- Tinbergen, N. 1963. On aims and methods of ethology. *Zeitschrift für Tierpsychologie* 20: 410–33.
- Tomasello, M. 1999. *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
- Townsend, D. J. & Bever, T. G. 2001. *Sentence comprehension: the integration of habits and rules*. Cambridge, MA: MIT Press.
- Trehub, S. E. 1976. The discrimination of foreign speech contrasts by infants and adults. *Child Development* 47: 466–72.
- Treves, A. 2005. Frontal latching networks: a possible neural basis for infinite recursion. *Cognitive Neuropsychology* 21: 276–91.
- Turing, A. M. 1936. On the computability of numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society, ser. 2* 42: 230–65.
- Uexküll, J. von 1921. *Umwelt and Innenwelt der Tiere*. Berlin.
- Uriagereka, J. 1998. *Rhyme and reason*. Cambridge, MA: MIT Press.
- Uriagereka, J. 2008. *Syntactic anchors*. Cambridge: Cambridge University Press.
- Uttal, W. R. 2001. *The new phrenology: the limits of localizing cognitive processes in the brain*. Cambridge, MA: MIT Press.
- Vargha-Khadem, F., Gadian, D. G., Copp, A., & Mishkin, M. 2005. FOXP2 and the neuroanatomy of speech and language. *Nature Reviews Neuroscience* 6: 131–8.
- Viertel, J. 1966. Concepts of language underlying the 18th century controversy about the origin of language. *Georgetown Monograph Series on Language and Linguistics* 19: 109–32.
- Villiers, J. G. de 2005. Can language acquisition give children a point of view? In J. Astington & J. Baird, eds., *Why language matters for theory of mind* (pp. 186–218). Oxford: Oxford University Press.
- Voltaire 1756. *Essai sur l'histoire générale et sur les moeurs et l'esprit des nations*. Paris.

- Vouloumanos, A. & Werker, J. F. 2004. Tuned to the signal: the privileged status of speech for young infants. *Developmental Science* 7: 270–6.
- Vouloumanos, A. & Werker, J. F. 2007. Listening to language at birth: evidence for a bias for speech in neonates. *Developmental Science* 10: 159–64.
- Wang, R. F., Hermer-Vasquez, L., & Spelke, E. S. 1999. Mechanisms of reorientation and object localization by children: a comparison with rats. *Behavioral Neuroscience* 113: 475–85.
- Watanabe, A. 2001. *Wh-in-situ languages*. In M. Baltin & C. Collins, eds., *The handbook of contemporary syntactic theory* (pp. 203–25). Oxford: Blackwell.
- Watson, J. B. 1919. *Psychology from the standpoint of a behaviorist*. Philadelphia: J. B. Lippincott.
- Watson, J. B. 1924. *Behaviorism*. New York: People's Institute Publishing Company.
- Weinberg, S. 1993. *Dreams of a final theory: the search for the fundamental laws of nature*. London: Radius.
- Weinreich, M. 1945. Der YIVO un di problemen fun undzer tsayt. *YIVO Bletter* 25: 3–18.
- Wernicke, C. 1874. *Der aphasische Symptomencomplex*. Breslau: Cohn & Weigert.
- West-Eberhard, M. J. 2003. *Developmental plasticity and evolution*. Oxford: Oxford University Press.
- Wexler, K. 1999. Maturation and growth of grammar. In W. C. Ritchie & T. K. Bhatia, eds., *The handbook of language acquisition* (pp. 55–109). San Diego: Academic Press.
- Wexler, K. & Culicover, P. 1980. *Formal principles of language acquisition*. Cambridge, MA: MIT Press.
- Wexler, K. & Manzini, R. 1987. Parameters and learnability in binding theory. In T. Roeper and E. Williams, eds., *Parameter setting* (pp. 41–76). Dordrecht: Reidel.
- Wheeler 2008. God's machines: Descartes and the mechanization of mind. In P. Husbands, O. Holland & M. Wheeler, eds., *The mechanical mind in history* (pp. 307–30). Cambridge, MA: MIT Press.
- Wilczek, F. & Devine, B. 1987. *Longing for the harmonies: themes and variations from modern physics*. New York: W. W. Norton.
- Wittgenstein, L. 1953. *Philosophical investigations*. Oxford: Blackwell.
- Wynn, K. 1990. Children's understanding of counting. *Cognition* 36: 155–93.
- Wynn, K. 1992. Children's acquisition of the number words and the counting system. *Cognitive Psychology* 24: 220–51.
- Xu, F. & Spelke, E. S. 2000. Large number discrimination in 6-month-old infants. *Cognition* 74: B1–B11.
- Yang, C. D. 2002. *Knowledge and learning in natural language*. Oxford: Oxford University Press.
- Yang, C. D. 2004. Universal grammar, statistics, or both. *Trends in Cognitive Sciences* 8: 451–6.
- Yang, C. D. 2005. On productivity. *Yearbook of Language Variation* 5: 333–70.
- Yang, C. D. 2006. *The infinite gift*. New York: Scribner's.
- Yolton, J. W. 1983. *Thinking matter: materialism in eighteenth-century Britain*. Minneapolis: University of Minnesota Press.
- Yoshida, M., Aoshima, S., & Phillips, C. 2004. Relative clause prediction in Japanese. Paper presented at the seventeenth Annual CUNY Sentence Processing Conference.
- Zurif, E. B. 1980. Language mechanisms: a neuropsychological perspective. *American Scientist* 68: 305–11.

Index

- acceptability, 134, 138–9
acceptability judgments, 11, 136, 138–9, 140
action, 122, 124
adaptation, 127, 164, 167, 181
adverb, 85–6
agent, 122, 124
agreement morphology, 66–9, *see also* inflection
algorithm, 31
Al-Sayyid Bedouin Sign Language, 44–5
amusia, 181
ant, *see* desert ant
ant lion, 26
aphasia, 152–4, 181, *see also* language disorder
approximate number system, 123–4, 125, 129–30, 165
Aristotle, 111
Arnaud, Antoine, 29
Arthur, Wallace, 88
Aslin, Richard N., 100–1
associationist, 100
astounding hypothesis, 149
auditory system, 154–5
- Babbage, Charles, 33
babbling, 51, 166, *see also* language production stages
Baker, Mark, 83, 89–90
- Basque, 92–3
bat, 171
Beauzée, Nicolas, 29
bee, 165–6
behaviorism, 15–20, 111, 133
Ben Shalom, Dorit, 159–60
Berkeley, George, 29
Berra, Yogi, 182
Bever, Thomas G., 137–8, 140
Bickerton, Derek, 44
binding, 32
binding problem, 161
biolinguistics, 4, 186
bird, 171, 178, *see also* finch, songbird
Bloom, Paul, 19, 97
Bouillaud, Jean-Baptiste, 151
brain, 149–62, 173
brain imaging, 140, 150–7
British Museum search, 39
Broca, Paul, 150–3, 157–8
Broca's aphasia, 153
Broca's area, 151, 153, 162, 180–1
Brodman, Korbinian, 151–2
Bulgarian, 89
- Carroll, Lewis, 58
Cartesian, 26–33, 118–19
 Cartesian linguistics, 26–33
 Cartesian naturalism, 29
cartography, 156, 158, *see also* spatial map

- cat, 48–9
- center embedding, 134, 142, *see also*
 - unbounded discontinuous
 - dependencies, rewrite rule (system)
- cetaceans, 171
- chaffinch, 22
- chimpanzee, 22, 171, 178
- Chinese, 38, 84–5, 89
- Chomsky, Noam, 3–5, 11, 15–20, 21,
 - 24–5, 26–32, 43, 46, 59–72, 83–4, 87,
 - 113, 133–7, 141, 167–72, 180, 183–8
- Christopher, 182
- chunking, 80
- Church, Alonzo, 33, 36
- Church-Turing thesis, 36
- cognitive impenetrability (of modules),
 - 127
- cognitive neuroscience, *see* neuroscience
- cognitive revolution, 13, 14, 57, 79–80, 89,
 - 163
- cognitive science, xi, 11, 36, 54, 177, 183–7
 - holy grail of, 111, 136, 149
- cognoscitive power, 31, 118
- cohesion, 122
- combination, 167
- comparative cognition, 121
- comparative genomics, 179
- comparative method
 - in linguistics, 90–3
 - in evolutionary studies, 163–72
- comparative psychology, 121
- comparative studies, 121, 165–72
- competence, 78–9, 133–6, 168
- complement, 72
- compositionality, 32
- computation, 31–3
 - limits of, 36
- computational acquisition, 98
- computational language learning, 99
- computational modeling, 98
- computational organology, 180–2
- computational theory of mind, 15, 31, 161,
 - 181, 183
- concatenation, 159
- concept, 29–30, 116–19, 131, 166–7, 170
- conservativity, 76–7
- constituent, 58, 71, 141
- contact, 123
- Context-free Phrase Structure Grammar,
 - 62–9
- continuity, 123
- Continuity Hypothesis, 106
- cooking, 120–1
- copy, 77
- Cordemoy, Géraud de, 30
- core knowledge systems, 121–5, 127, 131,
 - 167
- corpus, 135
- cotton-top tamarin, 170
- counting routine, 130
- Crain, Stephen, 96
- creative aspect of language use, 17, 28–9,
 - 36, 59, 119, 132, 184, 186
- creole, 43–5, 53
- creolization, 95, 100, *see also* language
 - creation
- critical period, 48–54, 166
- crossing dependencies, 65–9
- cross-serial discontinuous dependency, 66
- Cudworth, Ralph, 30, 31, 118–19
- cues, 98, 103–4
- Darwin, Charles, 4, 23, 31, 53, 87, 99, 120,
 - 150, 170, 184, 187
- Darwinism, 163, 180, 190
- data input, 97
- data intake, 97
- Davidson, Donald, 112
- Dawkins, Richard, 190
- dead reckoning, 124–5, 128
- Degree-0 restriction, 104
- Dehaene, Stanislas, 182
- delayed language production, 52
- Dennett, Dan, 4
- derivation, 62–5
- derivational theory of complexity, 137–8,
 - 140, 142
- Descartes, René, 5, 15, 26–31, 44, 100, 118,
 - 120, 124, 132, 184
- descriptive grammar, 9
- desert ant, 124–5
- developmental psychology, 121–5
- dialect, 82, 89
- Dijksterhuis, Eduard, 27

- discreteness, 57, 69, 161, 165
- discreteness problem, 161
- disorder, 173–4, *see also* language disorder
- displacement, 68–9, 73, 98
- dissociation, 173, 174–5
- distance effect, 123, 128
- dolphin, 171
- domain-general theories of learning, 24
- downward entailing, 116
- Dresher, Elan, 108
- Du Marsais, César Chesneau, 29
- Dwyer, Sue, 12

- Eibl-Eibesfeldt, Irenäus, 26
- Einstein, Albert, 13, 52, 79, 119
- E-language, 5, 45
- Eldredge, Niles, 184
- electroencephalography (EEG), 156
- electron, 71
- empiricism, 19, 28–9
- endocentricity, 71–4
- English, 46–7, 83, 84–7, 89, 91–3, 96, 103, 113
 - auxiliary system, 66–9
 - Middle, 95
 - Old, 95
 - subject-auxiliary inversion, 46–7, 105, 106–7
- English Platonist, 30
- entailment, 114–16
- ethology, 14, 20–5, 166, 187
- European starling, 169
- evolution, 163–72, 180
 - of language, 163–72
- evolutionary biology, 99, 165–72, *see also* comparative studies
- evolutionary psychology, 127
- exceptions, 102–3
- extension(al), 59, 160
- externalism, 112
- externalization, 120, 164, 167, 171

- Faculty of Language (in the) Broad (sense) (FLB), 9, 168
- Faculty of Language (in the) Narrow (sense) (FLN), 9, 168
- feature blindness, 175

- Feynman, Richard, 11, 83
- figure-ground distinction, 184
- finch, 22, 178–9
- finite-state automaton, 59–62, 169
- first cognitive revolution, 14
- Fitch, Tecumseh, 167–72, 177, 180, 187
- Fodor, Janet D., 104–5
- Fodor, Jerry A., 29, 36, 126–7, 137–8, 140
- folk biology, 122
- folk psychology, 122
- folk science, 122
- formal language, 112
- Foxp2, 176–9
- FOXP2, 174–9
- fractal, 69–71
- fractality, 69–71
- Frege, Gottlob, 111–12
- French, 85–6, 93
- frequency, 105
- functional magnetic resonance imaging (fMRI), 150, 156

- Galileo, 15, 87–8
- Gall, Franz Joseph, 126, 154–5
- Gallistel, Randy, 23–5, 37, 39
- Gambell, Timothy, 101
- games, 121
- garden path, 140–2
- Garrett, Merrill, 137–8, 140
- General Grammar, 29, *see also* Universal Grammar
- generative grammar, 136–7
- Genie, 51
- genome, 88, 173
- Gentner, Timothy, 169
- German, 107
- German romantics, *see* romantics
- Geschwind, Norman, 150–1
- Gestalt psychology, 75, 184, 191
- Gleitman, Lila, 19, 58
- Gödel, Kurt, 33
- Goethe, Johann Wolfgang von, 119
- Gopnik, Myrna, 175
- Grammatical Specific Language Impairment (G-SLI), 181
- grammaticality, 138–9
- grammaticality judgment, 138

- granularity mismatch problem, 159
 graylag, 21–2
 Greenberg, Joseph, 86

 habituation technique, 50
 Hale, Kenneth L., 52, 91–3
 Halle, Morris, 7, 61, 91
 Harvey, William, 23–4
 Hauser, Marc, 23, 131–2, 167–72, 177, 180, 187
 head parameter, 86–7, 96
 headedness, 71–4, 86–7
 Herbert of Cherbury, 30
 Herder, Johann Gottfried von, 30
 heterochrony, 88
 heterometry, 88
 heterotopy, 88
 heterotypy, 88
 Hickok, Greg, 151–2
 hierarchy, 62, 69
 Hilbert, David, 34, 194
 Hobbes, Thomas, 15, 28, 29, 31, 33
 home sign, 44
Homo combinans, 171
Homo sapiens, 82, 131, 171
 Hubel, David H., 48–9
 Human (language), 87
 human uniqueness, 120–1, 131–2
 Humboldt, Wilhelm von, 36, 44, 59, 118, 184
 Hume, David, 15, 27, 29, 190
 humor, 24
 hypothesis space, 97, 105, 107

 idealization, 77–9, 98, 133
 I-language, 5, 59, 82
 I-linguistics, 59, 82
 imprinting, 21–2
 infinity (of language use), 59–69, 77–9, 160, 166, 184
 Inflection (Infl), 85–6
 information, 37, 59, 97
 integration, 182
 intension(al), 59, 160–1
 interaction effect, 113, 134, 137
 internalism, 112, 118–19, 184
 intransitive verbs, 91–3

 Isabelle, 51–2
 I-semantics, 112
 Italian, 85, 107, 157–8

 Jabberwocky, 58, 158
 Jackendoff, Ray, xiii, 12, 184, 186–7
 Jacob, François, 88, 180
 James, William, 21
 Japanese, 50, 83, 86–7, 96, 103, 157

 Kanizsa triangle, 58
 Kant, Immanuel, 30, 118
 Kinzler, Katherine, 124, 125
 knowledge of language, 78, *see also* competence

 label, 71
 Lancelot, Claude, 29
 language acquisition, 42, 45, 83, 95–108
 variational model of, 99
 Language Acquisition Device (LAD), 21
 language change, 95, 96, 100, 105
 language creation, 43–5, 53, 95, *see also* creolization
 language disorder, 144, 152–4, 173
 language growth, 45, 107
 language production stages, 51
 language of thought, 131
 Larson, Richard, xiii
 Lashley, Karl, 19
 Lasnik, Howard, 61
 Lau, E., 160
 laws, 121
 learning bias, 97–108
 left hemisphere, 153–4, 162, 166
 Leibniz, Gottfried, 27, 29, 30, 33, 118, 124
 Lely, Heather van der, 181
 Lenneberg, Eric, 22, 47–8, 53, 176
 Lerdahl, Fred, 12, 186–7
 levels of representation, 58, 159
 levels of structure, 58, 159
 Lewontin, Richard, 172
 lexicalization, 79, 131–2, 167, 170
 lexicon, 160, 164
 Lichtheim, Ludwig, 150
 Lightfoot, David, 104

- linear order, 73–4
- linearization, 74, 159
- linking hypotheses, 150, 153, 159
- Llull, Ramon, 33
- locality, 75–6
- Locke, John, 29, 100
- long-term memory, 102
- loop, 60
- Lorenz, Konrad, 18, 21–2
- Lyell-Müller illusion, 20, 126

- magnetoencephalography (MEG), 156
- Marantz, Alec, 137, 145, 161
- Marcus, Gary, 32
- Marr, David, 40, 74, 79–81
 - Marr's three levels of analysis, 79–81
- massive modularity, 121
- master genes, 88
- mathematics, 48, 122, 125
- Mayr, Ernst, 99
- McCloskey, Jim, 53–4
- McGilvray, James, 118
- meaning, 111–19
- mechanization of mind, 36
- mental modules, 121–2, 126
- mental organ, 121–2
- mentalism, 18
- Merge, 79, 141, 159
- Mikhail, John, 12
- Miller, George, 79, 133–4, 137, 141
- mind–body problem, 32
- minimal pairs, 139
- modal, 66
- modularity, 121–2, 125–7, 179–81
- module, 36, 121–2, 125–7
- Mohawk, 93
- monkeys, 123, 166–7, *see also* chimpanzee,
 - vervet monkey
- Monod, Jacques, 88
- Montague, Richard, 112
- moral sense, 48, 94, 185–6
 - moral judgment, 40
 - moral organ, 23, 184
- Moro, Andrea, 57, 157–8, 161–2, 181–2
- morpheme, 57, 72
- morphology, 58
- motherese, 45

- multilingualism, 52–3, 84
- music, 40, 48, 94, 180–1, 185–7

- naïve physics, 122
- nativism, 18
- natural geometry, 122, 124–5, 128–9
- natural language semantics, *see* semantics
- natural philosopher, 26, 149–50
- natural selection, 99, 127, 164
- Navajo, 83, 92
- navigation, 122, 128–9
- negative evidence, 45
- negative facts, 9, 112, 114–16, 119, 135
- negative knowledge, 9
- negative polarity items, 114–16
- Neural Darwinism, 49
- neural network, 32, 38–9
- neuropsychology, 151
- neuroscience, 136, 149–62, 168, 182
- Nevins, Andrew, 38–9
- Newport, Elissa L., 100–1
- Newton, Isaac, 15, 18, 36, 57, 119, 141, 192
- Nicaraguan Sign Language, 44–5, 82
- number sense, 122, 123–4, 129–30, 182
- numerosity, 123–4, 129–30

- object mechanics, 122–3
- offline measurements, 140
- online measurements, 140
- ontological incommensurability problem,
 - 159
- organs, 24
- overgeneralization, 101–3, 106

- parallelism constraint, 7, 76
- parameter, 83–94, 95–7, 103–8, 186
- parsing, 104
- Patel, Aniruddh, 180
- path integration, 124
- Pavlov, Ivan, 16
- Pernau, Johann Ferdinand Adam von, 26
- Peirce, C. S., 38
- performance, 78–9, 133–6
- periscope, 71–4
- Phillips, Colin, 139, 140–5, 160, 174
- phonological bootstrapping, 101
- phonology, 58, 158, 159, 164

- phrase, 58, 71, 141
- phrase structure grammar, 62–9, 169
- phrase structure rule, 67
- phrenology, 126, 154–5, 158
- Piaget, Jean, 107
- Piattelli-Palmarini, Massimo, 20, 135, 171, 174–9
- pidgin, 43–4, 53, 107
- Pietroski, Paul, 113, 119
- Pinker, Steven, 13, 19
- pit viper, 25
- Pizlo, Zygmund, 74
- Plato, 31, 42, 107, 111, 124
- Poeppel, David, 13, 151–6, 158–61, 177, 180
- polysynthetic languages, 93
- population thinking, 99
- Portner, Paul, 116
- Port-Royal grammarians, 29
- positron emission tomography (PET), 150, 156
- Post, Emil, 33
- poverty of (the) stimulus, 42–7, 105, 185
- pragmatics, 58
- pre-existence, 31
- prescriptive grammar, xiii, 9
- Priestley, Joseph, 150
- probability, 17, 99
- productivity, 102
- projective geometry, 74
- psycholinguistics, 138
- puberty, 51–2, 83
- Pylyshyn, Zenon, 127
- quantification, 32
- question formation, 84–5, 89
- Quine, Willard v. O., 41
- rationalist/rationalism, 29, 32, 184
- rationalist psychology, 30, 31
- reaction time, 138, 140
- recursion, 32
- recursive, 63
- reduction, 182
- reduplication, 38
- reference, 112–13, 117–18, 134
- regularity, 101–3
- Reimarus, Hermann Samuel, 26
- reinforcement, 17
- replacement, 73
- representation, 29, 37, 100, 191
- representational theory of mind, 15
- representationalist, 100
- retinotopic map, 154
- retinotopy, 155
- rewrite rule (system), 62–9, 169
- Richards, Norvin, 91
- Rizzi, Luigi, 180
- romantics, 29, 30, 32
- Russian, 89
- Saffran, Jenny R., 100–1
- Saussure, Ferdinand de, 92
- savant, 182
- Schlegel, Karl Wilhelm Friedrich von, 29
- science-forming faculty, 119, 121
- Science of Man, 4, 13, 15, 190
- seal, 171
- second language learning, 52–3
- selection, *see* natural selection
- semantics, 58, 111–19, 121, 131, 134, 164
- sensitive period, 48–9, *see also* critical period
- Shakespeare, William, 119
- Shannon, Claude, 37–8
- sign (languages), 44–5, 51, 53
- signature limits (of core knowledge systems), 122, 125, 130
- size effect, 123
- Skinner, B. F., 15–20, 118
- Slavic languages, 89
- Smith, Neil, 182
- society finch, 22
- Socrates, 42, 46
- somatosensory homunculus, 155–6
- somatotopy, 155
- songbird, 166, 171, *see also* finch
- Spanish, 85
- spatial expressions, 128–9
- spatial map, 154–6
- Specific Language Impairment (SLI), 175, 181
- specificity, 179–81
- specifier, 72

- Spelke, Elizabeth, 120, 124, 125, 128, 130, 170
 Spelke-objects, 123
 spider, 26
 Spinoza, Baruch, 13, 27, 30
 spoonerism, 144
 starling, 169
 statistical learning, 99–100
 stress, 101
 structure-dependence, 46–7, 72
 subitizing, 123, 125, 129–30, 170
 successor function, 35
 syllable, 100–1
 syntactic bootstrapping, 58
 syntax, 58, 158, 159, 164, 177
 Szent-Györgi, Albert von, 79

 tail-embedding, 134
 tamarin, 170
 Tarski, Alfred, 111–12
 tense, 67–9
 theory of mind, 122, 131, 170
 Thornton, Rosalind, 96, 106
 Thorpe, W. H., 18
 Tinbergen, Niko, 18, 20, 21, 23, 163
 topographic map, 154–6, *see also* spatial map
 topology, 77
 tolerance principle, 98, 102
 tonotopic organization, 154
 tonotopy, 155
 trace, 77
 transcranial magnetic stimulation (TMS), 156
 transformation, 68–9
 transitive verbs, 91–3
 truth (-conditions), 112–13, 119, 134
 truth-values, 112
 Tsimpli, Ianthi-Maria, 182
 Turing, Alan, 33–6, 59, 79
 Turing machine, 33–6
 Turing test, 28
 type-token distinction, 32

 Uexküll, Jacob von, 21
 Umwelt, 22, 24, 54

 unbounded discontinuous dependencies, 63
 unification, 159, 182, 184
 uniformity, 69
 Universal Grammar (UG), 21, 29, 39, 44, 83–4, 91–4, 95–9, 106–8, 135, 157, 182
 updating mechanism, 97
 Uriagereka, Juan, 171, 174–9
 Uttal, William R., 154

 variable, 32, 160–1
 variational model of language acquisition, 99
 Vaucanson, Jacques de, 28
 vervet monkey, 165
 vision, 40, 48, 53, 74, 75, 94, 103, 154, 177–8
 visual cognition, *see* vision
 vocal imitation, 171

 walking, 54
 Wallace, Alfred Russell, 187
 wanna-contraction, 77
 Weber, Ernst Heinrich, 124, 202
 Weber's law, 123, 182
 Wernicke, Carl, 150–3
 Wernicke's aphasia, 153
 Wernicke's area, 151, 153, 162
 whale, 171
 Wheeler, Michael, 33
 Wiesel, Torsten N., 48–9
 Wilczek, Frank, 57
 Wittgenstein, Ludwig, 114
 word processing, 145
 word segmentation, 100–1
 words, 57, 72, 116, 166, 167, 170
 edge property of, 167
 working memory, 79–80

 X-bar schema, 72

 Yang, Charles, 99–101, 102–3

 zebra finch, 22, 178–9