CULTURAL SOFTWARE

A THEORY OF IDEOLOGY

J. M. BALKIN



Published with assistance from the Louis Stern Memorial Fund.

Copyright © 1998 by Yale University.

All rights reserved.

This book may not be reproduced, in whole or in part, including illustrations, in any form (beyond that copying permitted by Sections 107 and 108 of the U.S. Copyright Law and except by reviewers for the public press), without written permission from the publishers.

Set in Janson Roman by Rainsford Type, Danbury, Connecticut.

Printed in the United States of America

Library of Congress Cataloging-in-Publication Data

Balkin, J. M.

Cultural software: a theory of ideology / J.M. Balkin.

p. cm.

Includes bibliographical references (p.) and index.

ISBN 0-300-07288-0 (cloth: alk. paper)

ISBN 0-300-08450-1 (pbk. : alk. paper)

1. Ideology. 2. Culture. 3. Social values. 4. Justice.

I. Title.

B823.3.B25 1998

140—dc21 97-37011

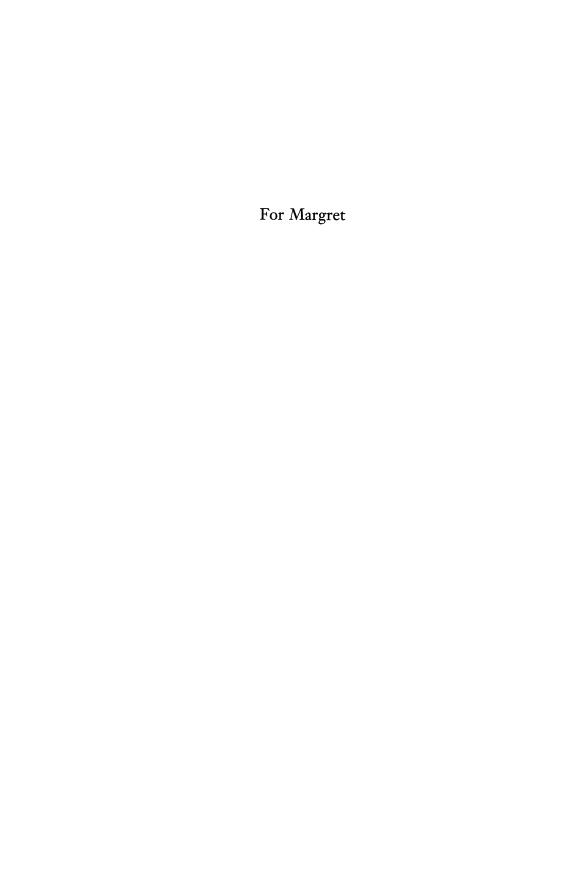
A catalogue record for this book is available from the British Library.

The paper in this book meets the guidelines for permanence and durability of the Committee on Production Guidelines for Book Longevity of the Council on Library Resources.

10 9 8 7 6 5 4 3 2

Cultural Software J. M. Balkin Yale University Press 1998 Online version

With the permission of Yale University Press, the author has made an online version of this book available under a Creative Commons Noncommercial Sharealike (by nc sa) license, accessible at the author's website at http://www.balkin.com.





Preface, ix

1 TOOLS OF UNDERSTANDING, 1

PART I CULTURE

- 2 BRICOLAGE AND THE CONSTRUCTION OF CULTURAL SOFTWARE, 23
- 3 MEMETIC EVOLUTION, 42
- 4 THE SPREAD OF CULTURAL SOFTWARE, 74

PART II IDEOLOGY

- 5 CONCEPTIONS OF IDEOLOGY, 101
- 6 AMBIVALENCE AND SELF-REFERENCE, 122
- 7 TRANSCENDENCE, 142

PART III METABRICOLAGE

- 8 CULTURAL HEURISTICS, 173
- 9 NARRATIVE EXPECTATIONS, 188
- 10 HOMOLOGIES AND ASSOCIATIONS, 216
- 11 METAPHOR, METONYMY, AND COGNITIVE MODELS, 242

PART IV UNDERSTANDING

- 12 THE POWER OF UNDERSTANDING, 261
- 13 KNOWLEDGE MADE FLESH, 286

Notes, 295

Index, 327



This is a book about culture and historical existence. It is a book about what cultures are and about what it means to be a person who lives at a certain point in history and whose individuality is shaped by a particular time and a particular culture. These questions have fascinated people for centuries; this book offers a distinctive answer to them. I argue that what unites the ideas of culture, personhood, and historical existence is information. To be part of a culture, to be socialized or acculturated, is to possess a certain kind of information—cultural know-how. Cultures are populations of individuals with relatively similar kinds of cultural information. To be a person is to be constituted by a particular kind of cultural information that exists at a particular point in time. The cultural information within human beings grows, changes, and evolves as we come in contact with others. It is reflected in our technology, in our institutions, and in the articulation of the values we hold most dear.

We are the bearers of this cultural information; indeed we are constituted by it. And its constitution of us is our constitution as historical beings. It is the source of our historical existence.

I call this cultural information *cultural software*. Because cultural software is the basis of all cultural understanding, it is the basis of the shared ways of understanding that people call ideology. The different beliefs and worldviews that human beings possess are the product of the evolution of cultural information that is instantiated in human beings and helps makes them the unique individuals they are.

People come to these kinds of issues in many different ways and from many different directions. I came to them through the study of law. It does not take long for lawyers to recognize that people's views about what the law is and

should be are often shaped by their beliefs about society and justice. Moreover, law is often asked to resolve disputes between people with very different views of the world. Because I was drawn repeatedly to the question of how ideology worked, I decided to write a book about the subject. But as with so many projects, this one took me in unexpected directions. What started out as a relatively straightfoward analysis of the nature of ideology ended up as a speculative essay about the philosophy of culture. Of course, the notion that the ideas in our heads take on a life of their own, and that ideas have us as much as we have them, is one of the major themes of this book.

In the last half of this century people have offered many theories that either assert or assume that individuals are socially and culturally constructed. These theories have tended to submerge the individual into the larger forces of society and culture. Perhaps partly in reaction to these trends, another group of theories has reasserted methodological individualism—the view that all social phenomena can be explained in terms of individuals, their actions and their mental states. Not surprisingly, these opposing approaches have symmetrical strengths and weaknesses, each explaining best what the other downplays or disregards.

I believe that a theory of culture must account for the uniqueness of each individual human being, while showing how the social and cultural forces shape us and produces our individuality. And I believe that individuality, like human freedom, is produced through culture, not in spite of it. The theory of cultural software offered in this book tries to explain why this is so.

Like methodological individualists, I would rather do without supraindividual entities. I do not think that these entities can adequately explain the production of shared beliefs or the presence of dissensus and disagreement within cultures and communities. Yet I also believe that a thoroughgoing methodological individualism is incomplete, for the individual is not the only unit of social explanation. Advocates of social construction have been looking in the wrong place; instead of looking above or beyond individuals, they should look deeper inside them. Instead of supraindividual entities, a theory of culture needs to take account of subindividual entities: the units of cultural transmission that help form individuals and create an economy of cultural development and exchange. These subindividual entities are cultural software; and they produce many of the effects that have led social theorists in the past to look outside the individual for explanations of the cultural.

This additional level of explanation makes it possible to view culture and society in a different light. We can see cultural software as something that both constitutes our interests and works against them. We can understand conventions and institutions as self-reproducing coordinated complexes of cultural software that have their own "interests" in survival and reproduction. And we can recognize how certain kinds of cultural software act like virtual parasites,

breeding unhappiness and injustice as they reproduce in human minds and institutions.

Postulating subindividual entities like cultural software also frees us from other forms of social explanation. We need no longer offer functionalist accounts to show why cultures and societies develop as they do. We need no longer assume that human conventions and institutions exist because they promote economic efficiency or human happiness, or because they solve important problems of social coordination. We need no longer offer "just-so" stories to explain away injustice or human misery as the necessary adjunct of achieving proper social functioning. Instead of functionalist accounts, we can offer evolutionary accounts, where the units of selection are not human beings but their cultural software, a cultural software that thrives and reproduces in the ecology of human minds. Nor need we explain human suffering as the necessary working out of a predetermined teleology of progress. Cultural evolution proceeds, but not toward any particular goal. We are its agents but not its puppets. We are its bearers but not its slaves.

Finally, the approach that I advocate in this book allows us to bring together many different and seemingly contradictory research programs in the human sciences. If human culture is bricolage, the solution of problems by adapting the various intellectual tools that lie to hand, the study of human culture—which is itself a part of culture—can hardly claim greater methodological purity. It should be a bricolage about that bricolage, a metabricolage that makes use of insights from many different sources and approaches. That is the entrepreneurial spirit of human reason. That is the spirit in which I undertook this volume.

When I told friends that I was writing a book entitled Cultural Software, some of them were amused. Others nodded knowingly and advised me that I am simply a victim of the age in which I live: an age in which new fortunes and new empires are being constructed daily out of computer programs and computer networks. After all, didn't philosophers impressed by Newtonian science offer mechanical models of human thought?

My response to my friends has been that we are always influenced by the age in which we live, we always absorb the intellectual tools that lie to hand. That is what it means to exist historically. That is one of the central themes of this book. Even if I deliberately eschewed the metaphor of software and hid my ideas beneath other conceptions less obviously inspired by the events of the late twentieth century, there would be influence enough. I use metaphors and ideas because they are useful, and in the hope that they will create a spark of recognition and excitement in others. It is true that no one would have compared cultural understanding to software two hundred years ago, but now that the comparison is possible, why not employ it?

Still others have objected to the metaphor of software on the grounds that computers are soulless, inhuman, and mechanical, and that human culture is too laden with value and emotion to be compared to information. But these concerns, too, strike me as misplaced. Cultural know-how is one of human-kind's most distinctive characteristics; we become who we are through the absorption and communication of information. And cultural information is not some inert form of data; it is skill, know-how, ability, empowerment. It does not do without human values but articulates them. It does not displace emotions but helps express them. It does not extinguish personhood but makes it complete. People who think that cultural information is soulless understand neither information nor the soul.

The metaphor of software, I predict, will in time seem no more forbidding than the metaphors of engines, or bicycles, or railroads seem today. Nowadays, some people think that steam-powered railroads are romantic. Yet for many in the nineteenth century, they were the very symbol of heartless mechanism. No doubt *software* will change its cultural connotations over time. Perhaps someday the word itself will seem hopelessly quaint and antiquated. I can only pray that the ideas contained in this book will not meet a similar fate. It is my hope that a few of these notions will grow, develop, flourish, and spread to other minds, helping those who absorb them to understand this world and themselves a little better. More than this an author cannot ask.

I was honored to be able to present the argument of this book as the 1997 Julius Rosenthal Lectures at Northwestern University Law School. I am grateful to the dean and to the faculty for their generous invitation and their warm hospitality.

Many friends and colleagues have commented on chapters of this book in its many previous drafts. I would like to thank in particular Bruce Ackerman, Tom Baldwin, Step Feldman, Owen Fiss, Liane Gabora, Bob Gordon, Susan James, Larry Lessig, Sanford Levinson, Jay Mootz, Tim O'Hagan, Richard Posner, Thomas Seung, Reva Siegel, Lea VanderVelde, and R. George Wright for their comments, as well as participants at a workshop at the University of Chicago Law School where drafts of what became the first four chapters were presented. I am particularly grateful to an anonymous reader for Yale University Press whose trenchant criticisms greatly improved the book. An early version of Chapter 1 was published as "Ideology as Cultural Software," 16 Cardozo Law Review 1221 (1995). Portions of Chapter 7 were originally published in "Transcendental Deconstruction, Transcendent Justice," 92 Michigan Law Review 1131 (1994).

Finally, I want to thank my wife, Margret Wolfe, for her faith in me, her emotional support, and most of all, her love. This book is dedicated to her.



In the Pirke Avot, or "Ethics of the Fathers," a popular tractate of the Talmud, one finds a list of objects that God brought into being as the sun set on the sixth and final day of creation. These include the modus operandi for many of the miracles that the Lord would later display before humankind—for example, the well of the Earth that swallowed Korah and his followers, and the mouth of the ass that spoke to Baalam in the desert. It was as if, as sunset approached, and God knew that He was finished with the task of creation, He endeavored to bring into existence everything that might be needed later on. At the end of this list of necessary items is a curious addition: the first set of tongs, for, as the Talmud tells us, tongs can only be made with other tongs.

The idea of a tool that can be made only from another tool, and that is itself a toolmaking tool, lies at the heart of this book. My immediate concern is the phenomenon that theorists have called ideology, but my larger subject is human cultural understanding. I believe that the study of ideology must be dissolved into this larger concern. We must break down what previous thinkers have called ideology into distinct and analyzable mechanisms. We must replace the study of ideology with the study of diverse ideological effects produced by human thought, effects that together produce the phenomenon called ideology. At the same time, we must expand the concept of ideology by absorbing it into the more general study of cultural understanding. So we must proceed in a dual movement: dividing ideology into its variegated mechanisms, and viewing these mechanisms as special cases of the ordinary processes and operations of human thought.

The metaphor of the toolmaking tool unites these two gestures. The study of ideology is the study of tools of human understanding produced in, by, and through human culture. It is the study of the cumulative creation of these tools through the use of previously existing tools of understanding, and the study of the consequences of this recursive manufacture. To understand ideology we must understand the tools of human understanding—with respect to both their advantages and their deficiencies, their intended and their unintended consequences, their ability to empower us and their ability to exercise power over us.

Ideology is a much-contested term these days.² Some social theorists think that it has outlived its usefulness, particularly given its historical connections to the Marxist tradition and Marxism's many internal disputes. They prefer instead to talk about discourse, episteme, habitus, tradition, language game, interpretive community, and a host of other terms for characterizing the social nature of human thought. Each of these terms has a slightly different meaning. Each justifies its particular stance by a different theory. Yet each points at the same basic set of issues—the socially generated and socially sustained ways in which human beings understand and constitute their world. And regardless of the particular terminology used, each of these approaches produces different versions of a theory of ideology. When I speak of "the theory of ideology" in this book, I refer to their collective concerns. As I shall stress repeatedly, the distinctive problems faced by a theory of ideology do not vanish when we change our focus to concepts like discourse.

Ideology and the Philosophy of Culture

The theory of ideology, like the study of discourse associated with postmodernism, has always been part of a larger endeavor—the philosophy of culture. The ancient Greeks distinguished between *physis*, the world of nature, and *nomos*, the world of convention, law, and culture. The philosophical study of nomos includes ethics and political theory. But it also includes culture itself as a philosophical problem and an object of study.

The philosophy of culture has a rich tradition and many illustrious fore-bears, of whom Vico, Rousseau, Kant, and Hegel are perhaps the most prominent. Some of the questions the philosophy of culture asks are these: What is the relationship of culture to human existence and human history? What role does culture play in producing the faculty of human reason? Is human history, and hence the history of culture, a tragedy or a comedy, or is it a story with no determinate end and no narrative coherence or unity? Later philosophers, like Marx and Nietzsche, emphasized a further question: the question of power. What power do culture and cultural forms have over individuals? How can individuals recognize this power, and what, if anything, can or should they do about it? The study of discourse, like the study of ideology, is merely the latest

in a series of approaches to the philosophy of culture. The basic questions it asks are very much the same, and the problems it encounters are very similar.

Although people use the term *ideology* in many different ways, they are usually invoking one of two basic conceptions. The first sees ideology as a worldview, an intellectual framework, a way of talking, or a set of beliefs that helps constitute the way people experience the world. In this conception, ideology is a relatively neutral term. The second conception of ideology is distinctly pejorative. Ideology is a kind of mystification that serves class interests, promotes a false view of social relations, or produces injustice. Alternatively, ideology is a way of thinking and talking that helps constitute and sustain illegitimate and unacknowledged relationships of power. It is a form of discourse in which oppressive power finds its home.³

This book offers a third position. Both the neutral and the pejorative conceptions of ideology describe different aspects of a deeper phenomenon. They describe effects produced by the tools of human cultural understanding. I call these tools of understanding cultural software. Hence my theory of ideology is a theory of cultural software and its effects.

Oppressive discourses, worldviews, belief structures, and mystifications all arise from the diverse tools of human understanding. The components of cultural understanding include beliefs and judgments. But they also include cognitive mechanisms that help produce and fashion beliefs and judgments. These cognitive mechanisms include, among other things, heuristics for decision, narrative structures and social scripts, conceptual homologies (A is to B as C is to D), metaphor and metonymy, and methods of ego defense. Each of these cognitive mechanisms can be beneficial and useful in certain contexts, but in others each can mislead and help produce or sustain unjust conditions. The tools of human thought are both helpful and hurtful, depending upon when and how they are used. Recognition of the simultaneous advantages and disadvantages of our tools of understanding—the inevitable connections between heuristics and their limitations, between blindness and insight—is central to the argument of this book. I call this the ambivalent conception of ideology.

This approach replaces both the neutral and pejorative conceptions of ideology. As in the neutral conception, we still study how worldviews and systems of discourse are produced. But we do this by investigating the diverse mechanisms of cultural understanding, and we do not take a neutral or detached view toward their effects. Our tools of understanding simultaneously enable and limit our understanding, empower us and have power over us. When our cultural software helps create or sustain unjust conditions, I say that it has *ideological effects*. But our tools of understanding do not always produce these effects. Hence ideology, in the pejorative sense, is not a phenomenon separate from the general mechanisms of cultural understanding; it is an effect produced

by these mechanisms when they are placed in particular contexts and situations. I retain the familiar adjective ideological to describe these contextually produced effects.

The metaphor of "cultural software" proposes that we can compare certain features of culture, and of the way that culture operates, to the software that is installed on a computer and that allows a computer to process information. Simply put, cultural software enables and limits understanding as software enables and limits a computer. Although this can be a helpful metaphor, it can also be misunderstood. With this in mind, I want to discourage two likely misinterpretations.

First, I do not believe that the human mind works like any existing computer. Nor do I believe that thinking is primarily a mechanical or algorithmic process. On the contrary, I shall insist throughout my argument that human thinking is distinguished by its symbolic and metaphoric character and by its fundamental motivation in human values.

The growth of cognitive science and the search for forms of artificial intelligence have led naturally to comparisons between human beings and computers. One of the most important debates currently raging in the philosophy of mind is the extent to which mind should be defined functionally in terms of information states, like those in a computer. Some philosophers of mind have gone so far as to argue that the human mind is essentially indistinguishable from a computer, while others have asserted that the intentional nature of human intelligence makes such comparisons thoroughly inappropriate.4

Although these debates are interesting, they are to a large extent peripheral to the concerns of this book. Unlike most cognitive scientists and most philosophers of mind, I focus on the mind's relationship to culture and not its ultimate structure.5 Howard Gardner has noted that although most cognitive scientists "do not necessarily bear any animus against the affective realm, against the context that surrounds any action or thought, or against any historical or cultural analyses, they attempt to factor out these elements to the maximum extent possible."6 It is quite possible that the computational metaphor of mind has encouraged this trend. For these reasons, Jerome Bruner, himself one of the founders of the cognitive revolution, has recently called for a renewed emphasis on "the concept of meaning and the processes by which meanings are created and negotiated in a community." These concerns lie at the heart of this book; they motivate my use of the idea of cultural software. I use this metaphor to illuminate the ways in which human beings are constituted by and express their shared values within a culture.

Second, the idea of cultural software suggests an opposition to "biological hardware." But we cannot distinguish between "hardware" and "software" in humans in the way we can for computers. Each individual has a unique brain structure that is not merely the product of genetic inheritance but is shaped and organized in part by her experiences and activities, especially those in early childhood. As we are programmed through social learning, our physical brain structure is also changing. And the consequences of human beliefs and cultural activities affect human populations and human bodies in countless ways. Thus, it is highly misleading to think of individuals as consisting of identical hardware into which identical copies of software are installed.⁸

The idea of cultural software is not designed to suggest or defend a neat division between the cultural and the natural. Rather, it directs our attention to the know-how that is part of every human being and that is shared by and transmitted between human beings through communication and social learning. This know-how is our cultural software. The ability to communicate and engage in social learning and thus pass on cultural know-how is an essential aspect of our nature as human beings. The most remarkable result of human evolution is that it is in our nature to be cultural. We are by nature cultural creatures.

To imagine ourselves as cultural creatures is not to imagine ourselves as infinitely malleable; this assumption misunderstands the ways in which culture grows out of nature. The instincts and motivations that we have inherited from our genes are not abandoned or displaced by social learning. They are refined and articulated, distorted and exaggerated, extended and supplemented by experience and social learning. What is made is always made from materials already given, and its character and its limitations are shaped by those materials. In such a way the present is always made from the past. We can fashion a purse from a sow's ear, but it will be the kind of purse that can be so fashioned.

Yet at the same time, culture has a cumulative power. The present can only be made from the past, but the future can be made from the present. And so as culture is transmitted and transformed, it opens up ever new horizons of human possibility. In Chapters 2, 3, and 4, I shall argue that just as our biology has evolved through transmission of genes, our cultural software is also evolving through cultural transmission—although in importantly different ways and at much faster rates. These genetic and cultural processes necessarily interact with each other; this interaction is only one of the many ways in which the cultural forms part of and is continuous with the natural world.⁹

Each human being is born with the ability to absorb and communicate previously developed culture—to possess cultural software and transmit it to others. Because we can make culture part of us, we can be the living embodiments of previous cultural development, just as we can be the embodiments of previous genetic evolution. And because we can do this, we are also historical beings. We can absorb, reflect, and transmit the cultural know-how available at our particular moment in history. We can be the carriers of a certain kind

of cultural software, the kind present at our particular moment in history, and we can be the vehicle for its transformation into the cultural software that will be absorbed by future generations. We can, and indeed we must, stand in complicated lines of inheritance and innovation. To be the bearer of a particular kind of cultural software, a configuration existing at this time and at no other, is what it means to be a historical being, to exist historically.

History in this sense is a peculiarly human phenomenon; the Grand Canyon changes over time, but only human beings have history. Or more accurately, the Grand Canyon has a natural history, but only human beings have a cultural history, which is history proper. Human beings begin to have history only at the moment when they enter into culture, which is also the moment that they begin to create collectively shared and created tools for understanding the world and articulating their values.¹⁰

Culture and cultural software are just such tools. They are tools used to make other tools. This has always seemed to me the deeper meaning of the Talmudic story; for when God created human beings on the sixth day of creation, one of His final acts was to bequeath to them a toolmaking tool, which is human culture.

For simplicity's sake, we might distinguish three kinds of cultural tools that human beings use, keeping in mind that they are difficult to separate in practice. (Moreover, this list is not intended to be exhaustive.) The first is technology, the second is institutions, and the third is cultural know-how, or what I call cultural software. It consists of the abilities, associations, heuristics, metaphors, narratives, and capacities that we employ in understanding and evaluating the social world. An example of technology is a computer. An example of an institution is a bank. Examples of cultural software are knowing how to operate a computer, being able to dance the waltz, or being fluent in a particular language. Technology makes tools from materials, institutions make tools from human sociability, and cultural software makes tools from human understanding.

Although I have distinguished them analytically, in practice these three types of cultural tools are necessarily interdependent and interrelated. The institution of a bank, for example, may presuppose technology in the form of buildings, computers, furniture, and a workforce trained in a certain way, with certain understandings and abilities. The operation of technology requires know-how, and, conversely, certain skills and knowledges often presuppose certain technologies (as well as institutions already in place). Nevertheless, different philosophers of culture have emphasized some types of cultural tools more than others. For example, Marx emphasized the role of technology, and Vico emphasized the role of institutions. But the third type of tool—cultural software—is equally important. Without cultural software, our technology lies on the ground, rusted from dis-

use, and our institutions fall apart. The biblical story of the Tower of Babel is a good example of what becomes of technology and institutions without cultural software. Indeed, without cultural software, social institutions not only cannot be maintained; they cannot even get started.

Why Software? The Problem of Shared Understandings

The motivation behind the concept of cultural software is not the familiar desire to model the operations of the human brain on those of the digital computer. Rather, the point of this metaphor is to address and resolve a recurring problem in theories of cultural understanding: to explain how shared cultural understandings can be shared while still accounting for the considerable differentiation and disagreement in belief among members of the same culture or interpretive community.

To show how this problem arises, I shall use as an example Hans-Georg Gadamer's theory of cultural understanding. Gadamer argues that human cultural understanding is made possible by our location in a historically generated tradition. His theory is especially attractive because it draws an important connection between historical existence and cultural understanding. Gadamer emphasizes that human existence is existence in history; to be human means to exist in a historical tradition and hence to understand within and by means of this tradition.

Because we exist in a tradition, Gadamer claims, we bring certain prejudices or prejudgments to all of our understanding. But these prejudgments, far from being hindrances to our understanding, are in fact the preconditions of our understanding. They enable us to understand not only others within our own culture but people in other cultures as well. Thus, Gadamer asserts, "Understanding always implies a pre-understanding which is in turn pre-figured by the determinate tradition in which the interpreter lives and which shapes his prejudices."¹³

Gadamer does not view his theory of tradition as a theory of ideology; nevertheless, it provides an excellent starting point for my claim that ideology is a special case of ordinary cultural understanding. We might think of the ideology of Americans, for example, as a cultural tradition that shapes, directs, and facilitates their understanding. The prejudices and prejudgments associated with this tradition color Americans' views of the world and produce a distinctive take on various political questions. Indeed, we might be tempted to substitute the word *tradition* directly for the word *ideology*. We need only modify Gadamer's comparatively rosy view of the effects of prejudgments and prejudices on the understanding by emphasizing that these prejudices and prejudgments can as easily mislead as facilitate social understanding.

The close connections between ideology and the Gadamerian concept of tradition suggest the irony of the well-known critical exchange between Gadamer and Jürgen Habermas. Habermas worries that Gadamer's theory of understanding does not sufficiently take into account the distorting effects traditionally associated with ideology, when in fact Gadamer's concept of tradition can easily be adapted to provide an account of how ideological thinking occurs. Conversely, Gadamer insists that if successful understanding ever occurs, it must occur through a historically generated tradition with its prejudgments and prejudices; yet through this argument, Gadamer simultaneously demonstrates the inevitability of ideological limitations on thought. In short, although Gadamer's account of cultural understanding was designed to show how understanding can succeed, it also provides an account of how cultural understanding can go wrong.

In spite of its considerable utility for a theory of ideology, Gadamer's theory of cultural understanding creates a series of puzzles. First, it is not clear what kind of entity a tradition is and how it is possible to live in it. Where does the tradition exist so that we are able to live in it? If we live together in a tradition, it is surely not in the same way that two people live together in a house. Moreover, Gadamer wants simultaneously to insist that the tradition we live in also inheres in us, so that we are both inside it and it is inside us. In the alternative, one might say that we share in a tradition; but do we share it like a piece of clothing (which only one person can wear at a time), like a pie (from which we take separate slices), like a parent (having a common causal origin), or like an experience (having been subjected to roughly the same causal forces)? Finally, even though we are inside the tradition and it is simultaneously inside us, the tradition continues to exist after we (or any other individual) leave the community or die. Yet people are somehow also able to bring their traditions with them to new places after they leave their communities. To make sense of these puzzles we must know what kind of object a tradition is, where it may be found, and where, if anywhere, it continues to exist after individuals no longer form a part of it. Stephen Turner has called these various kinds of difficulties the problem of location.¹⁵ Note, however, that the word *location* is appropriate only because the metaphors generally used to describe tradition are spatial in nature: we say that we live in a tradition, the tradition inheres in us, certain behaviors or persons are outside of the tradition, and so on. In fact, the problem of location is really a problem of ontology: it is the question of what kind of object or entity a tradition is, given that we use these spatial metaphors to describe it.

Second, Gadamer says that the tradition is responsible for people having the kinds of prejudices and prejudgments they have. By implanting these prejudices and prejudgments, the tradition facilitates and empowers our understanding.¹⁶ But Gadamer does not tell us exactly how the tradition does this.

Gadamer believes that tradition is disseminated through communication and language.¹⁷ But that is only half an answer. What exactly is the thing that is disseminated, and how does it have causal effects on human intelligence? Turner calls this the problem of transmission.¹⁸ Once again, however, this expression reflects the standard metaphors employed: We hand over traditions, we implant them in others, we transmit them. In more general terms, it is the problem of *causation*—we need to know what kind of causal nexus exists between the tradition and individual human intelligence and/or behavior. Moreover, the question of causation is also the question of power, for it is the question of how traditions can have power over individual minds. Indeed, a recurring problem in theories of ideology has been some version of this question—the question of how ideas can have power over people.

Third, if tradition inheres within each individual in a culture and shapes each individual's apparatus of understanding, why do individuals ever differ in their understandings of the same tradition? Why, for example, do American constitutional lawyers disagree about the meaning of the Constitution if all of them are part of the same constitutional tradition? As Gadamer himself recognizes, one of the most interesting features of a cultural tradition is that its content and scope are always being tested and contested by the individuals who live within it.¹⁹ Yet how is this possible if all share in the tradition equally, or if the tradition inheres in each individual in the same way? This is the problem of differentiation. It is the flip side of the problem of transmission or causation, and, not surprisingly, it is sometimes neglected in theories that are trying to show how shared cultural understandings are shared. Ironically, one of the hidden dangers that any explanation of shared cultural understandings faces is that it will prove too much—that it will explain more uniformity of thought, belief, and action than actually exists in a given culture. Such accounts suppress the heterogeneity and dissensus that exist among the cultural understandings of any group of individuals. To be successful, then, a theory of shared cultural understanding must show not only why understandings are shared, but also why they are not shared—why no two people view the cultural world in exactly the same way, and why in any culture there are always mistakes, misunderstandings, and disagreements. The standard response that differentiation occurs because the tradition has unclear boundaries does not solve the problem but simply returns us to the spatial metaphor (a tradition has boundaries like a country). Thus, it raises anew the questions of location and transmission, or ontology and causation—what kind of thing could a tradition be for it to have unclear boundaries, where is this thing located, and how is it transmitted into each individual mind?

The problems of causation and differentiation, in turn, are related to a final problem, the problem of change or *transformation*. The tradition we live under

today is not the same in all respects as the tradition that existed two hundred years ago. The prejudices and prejudgments of one generation are often different from those of their children or grandchildren. Yet in spite of these changes, the tradition continues to be shared, although the content of what is shared has become different. How does this change occur, what produces it, and how is widespread agreement among members of the community preserved during this process? Solutions to these problems often raise the problems of ontology, causation, and differentiation in new guises: For example, if change occurs because of defects in transmission of the tradition to individuals, the transmission must be defective for all members in the same way if agreement is to be preserved on the terms of the newly changed tradition. If change occurs through individual differentiation, we need to know how agreement between individuals was ever maintained in the first place and how it is now obtained on new grounds. So the problem of transformation brings us back to the same old puzzles—what kind of thing is a tradition that it can change or be changed in this way, what kind of causal efficacy does it have over individuals, how is it implanted in them, and how is it implanted in the same way?

Although I have used Gadamer's theory of tradition to discuss these problems, they arise for many different types of entities and many different kinds of social theories that purport to explain the existence and effects of shared social understandings. If we were to substitute for Gadamer's "tradition" the idea of a "collective consciousness," an "Objective Spirit," a "habitus," a "practice," an "episteme," an "interpretive community," or a "form of life," the same questions of ontology, causation, differentiation, and transformation would arise again, albeit in slightly different ways. Gadamer's theory of tradition is one in a long line of approaches designed to show how shared understandings are shared. For convenience, we may group these approaches into three basic types, which I call the supraindividual, the behavioral, and the Kantian approaches.

The first type of solution, of which Gadamer's appears to be an example, postulates a supraindividual entity that somehow does the work of regulating or ordering the minds of individuals. Examples would include Hegel's notion of an Objective Spirit or Durkheim's notion of a collective consciousness. In these theories, a single entity existing over and above individual minds guarantees the shared nature of cultural understandings. This entity may be a supraindividual consciousness or, in the case of Gadamer's tradition, an entity whose nature is largely unelaborated. Not surprisingly, such theories create puzzles about what the supraindividual entity is, where it resides, how it is shared by individuals, what force it has over individual minds, and how disagreement and disputes are possible.

A second solution turns instead to behavior. It asserts that shared practices

of understanding are explained by the existence of shared conventions of social behavior. The usefulness of this solution depends on what sorts of things these conventions are. If conventions are viewed as entities that exist over and above individual minds, for example, they threaten to become just another version of an Objective Spirit, a collective consciousness, or a tradition, and they face similar difficulties.

We might try to avoid these problems by asserting that conventions are agreements to behave in similar ways. Nevertheless, they cannot be conscious agreements, for most people never consciously decide to adopt them. Alternatively, we might insist that by conventions we mean nothing more than regularities of behavior. Unfortunately, this solution leads to a problem of circularity. It is precisely these regularities of behavior that a theory of shared cultural understanding hopes to explain. The claim that shared understandings are shared by virtue of social conventions explains nothing. Nor does this solution explain how disagreement and differentiation are produced within conventions, for by definition such disagreements can occur only outside of them, or at those places or in those situations where social conventions run out.

Finally, an explanation of shared conventions in terms of similarities of behavior shifts our attention away from cognitive processes of meaning and understanding that must form part of each individual's conceptual apparatus. When we say that participants agree, we have not yet explained how they agree. Of course, it is an advantage of behavioral accounts that they avoid questions about what the internal mechanisms of cultural understanding are. Yet this advantage is also a disadvantage, for cognitive processes of meaning and understanding surely must be involved in the creation and maintenance of shared conventions. So the problem with this kind of account is that the solution it offers is just too easy-it simply declares victory and goes home without addressing the most difficult questions of how cultural understanding is regulated, transmitted, and maintained.

Some philosophers have tried to explain conventions in terms of similar or interlocking expectations. But these accounts cannot be purely behavioral; to explain shared understandings they must smuggle in the very sorts of concepts that raise the problems I have noted above. David Lewis, for example, defines conventions as regularities of behavior; yet his account depends on prior concepts like "common knowledge" of a state of affairs, mutual expectations, and individuals conforming to a regularity. "Common knowledge," in turn, depends on certain states of affairs indicating the same thing to everyone in a population.²⁰ The hermeneutical problems that I am concerned with enter at precisely these points in his account. Moreover, Lewis's account assumes that conventions solve problems of coordination based on people's preferences. But not all of the various types of shared meanings and beliefs that occur in a

culture can be explained as solving problems of coordination.²¹ In short, conventions do not explain shared understandings; they presuppose them.

A third type of solution to the problem of shared understandings is Kantian in spirit. According to this account, individuals within a culture understand the world in the same way not because of the existence of a supraindividual entity that regulates agreement but because each possesses an identical conceptual apparatus. Each individual's mind is similarly constructed and employs identical principles of conceptual construction, organization, and association. This common perceptual and cognitive apparatus produces and guarantees shared understandings.

I call this approach Kantian because it postulates something akin to Kant's notion of a transcendental subject. Although this transcendental subject is spoken of in the singular, it should not be confused with a supraindividual entity. It refers to the common features found in the subjectivity of all rational beings by virtue of their being rational.²² Its grammar is similar to that in the expression "the human eye," which refers to general features found in all normal examples of this organ. Similarly, talk about "the" transcendental subject refers to identical copies of a basic conceptual apparatus.

Kant used the concept of the transcendental subject to explain our understanding of very basic aspects of the natural world, but his idea can be extended much further. For example, Edmund Husserl argued that the transcendental ego gave each person the ability to comprehend eidetic essences and to perceive the world in terms of categories. One can also recognize a similar motivation in Claude Lévi-Strauss's concept of universal structures of the human unconscious that underlie all myths, or in Noam Chomsky's theory of a universal grammar that underlies all human language.²³

The most serious problems with this sort of approach occur in accounting for the differentiation of individual understandings and the transformation of shared understandings over time. If we limit our focus to explaining our common understanding of space and time in the natural world, these issues do not arise so urgently. However, we are trying to explain how people within a particular culture at a particular time share understandings that are partly different from those held by persons in other cultures and times. Thus we need a sort of "historicized transcendental subject," a common hermeneutic apparatus that is similar for all members within a culture but differs for people in different cultures and times. Yet in some sense this expression is a contradiction in terms; for what makes the transcendental subject transcendental is precisely its universality and resistance to historical variation.

One might attempt to avoid historicizing the transcendental subject by arguing for the universality of a limited set of features of human thought that explain some but not all features and varieties of shared cultural understandings.

In particular, one might retain a Kantian-style explanation of formal features of human cultural understanding, while conceding that more substantive features are subject to historical variation. For example, Lévi-Strauss argued that although different cultures have different myths, the basic principle of conceptual opposition is the same for all. Chomsky holds that although natural languages differ in many respects, they all share basic grammatical features. Unfortunately, this strategy leaves the basic problem unsolved. For we want to explain how shared understandings and substantive agreements occur within various cultures even though the nature and content of these shared understandings and substantive agreements differ in different cultures.

The great strength of Kantian-style explanations turns out to be their greatest weakness—they can guarantee shared understandings only so long as those understandings are unaffected by historical change. However, once we concede that each culture and time has its own version of a "transcendental subject"—a common conceptual apparatus that guarantees shared cultural understandings within its boundaries but that changes over time—we immediately face the familiar problems of transmission, differentiation, and transformation. Once history intrudes, we must explain what mechanism guarantees that people within particular cultures have roughly the same apparatus of understanding over time, and what causes this common apparatus to change in more or less identical ways for each person in the culture. The great advantage of supraindividual and behavioral accounts is that they can be historical in a way that a Kantian solution cannot. Yet they have their own difficulties in explaining the similarities (as well as the differences) among individuals' understandings.

The theory that I propose seeks to explain what people have traditionally called ideology as a special case of shared cultural understanding. But as our discussion has shown, the concept of shared cultural understanding itself needs serious explication. To describe the phenomenon of ideology, we need something like Gadamer's concept of tradition, but we must alter it considerably to avoid the puzzles that this and similar concepts produce. We need a way of explaining shared cultural understandings that avoids the defects of the three approaches mentioned above while combining their advantages. In short, we need something

- 1. that exists in each individual;
- 2. that shapes and enables individual understanding and cultural know-how;
- 3. that guarantees similarity of cultural understanding and know-how while permitting some variation, disagreement, and mistake among individuals within the same culture;
- 4. that changes and develops over time; and
- 5. that constitutes individuals as persons living in a particular culture at a particular point in history.

The best way to describe this thing is as a kind of cultural software. A copy of this software forms part of each person. Cultural software performs a function similar to Gadamer's "tradition": it provides us with the tools and preunderstandings that enable us to make judgments about the social world. Moreover, to the extent that people possess roughly similar copies, their cultural understandings are shared understandings. However, the theory of cultural software posses a different answer to the problem of location. While it is not clear where a Gadamerian tradition resides, our cultural software resides in us, because it is literally part of us.²⁴

Our cultural software is written and rewritten through social interaction and communication. These acts result in an economy of similarity and difference between the cultural software of different persons. This economy, in turn, produces convergence in cultural understandings as well as individual differentiation. It ensures that our cultural software is roughly similar to that of others in our culture even if it is by no means identical in all respects. Moreover, through this economy, the cultural software of individual human beings evolves over time. Our cultural software bears the marks and effects of previous development. It is the historical component of our human existence.

Thus what Gadamer calls tradition is not something that controls individual understanding but an effect produced by the cultural software of many individuals who have communicated with and thus affected each other's cultural software over many years. Saying that we "live in" a tradition means that we participate in an economy of cultural communication with others who have (or have had) roughly similar cultural software. Shared understandings are produced by the rough similarity of our cultural software, and regulated by our communication with others. Thus when we speak of "our cultural software" we do not refer to any supraindividual entity or collective consciousness. Rather, we mean only the collection of partly similar and partly different individual copies. In this sense, cultural software is the historicized analogue of the Kantian transcendental subject—it is a conceptual apparatus within each individual that prefigures cultural understanding but that can also change and evolve over time.

Cultural Software and the Construction of Persons

Behind this explanation of shared understandings lies a further and deeper intuition that motivates the metaphor of cultural software. Human beings are made of knowledge; we are the living embodiments of information. Everyone knows that human beings store information in their genes. And many people also know that human beings store information in their immune systems. This information, produced and shaped by the body's previous encounters with mi-

croorganisms, helps it ward off future disease. But human beings also embody a third kind of information—cultural information. We know things and we know how to do things because we live in cultures, and this ability, this knowledge, is central to our existence as persons.

Our human existence as embodiments of information, as bearers of cultural know-how, is the most basic motivation for the metaphor of cultural software. The comparison between cultural software and computer software encompasses two further ideas. The first is that software is an indispensable tool for processing information and performing tasks. The second is that software is an indispensable part of what we mean by "the computer."

Let me address these two points in turn. First, a computer needs software to process information. Without this software it cannot do its job; it cannot interact with the environment around it. If one boots up a computer without software, it just sits there and does nothing. One can type on the keyboard endlessly, but the computer will not respond, or at best it will spit out an error message. It cannot process information because it has nothing to process information with. At best its ability to process information is primitive and unhelpful. Only when we install software can it do anything useful, and even then the type of information it can process depends on the kind of software installed on it. The most massive supercomputer, installed only with a checkers program, still can only play checkers—though it can probably play checkers very quickly indeed. The potential power of the computer remains great, but its practical power is severely limited. As the power of the software grows and develops, so too does the practical power of the computer. In this way the potential abilities of the hardware are fully realized only through the development of increasingly elaborate software. Thus we might say in a very loose sense that software empowers hardware.

The second point is that this software is, to a very important extent, constitutive of the computer, or rather, what we unthinkingly call "the computer." Often what we mean by "the computer" is really the software together with the hardware. So I say that I wrote this chapter on my computer, but technically I wrote it using a word-processing program installed on my hard drive. For most of us, then, what we mean by "the computer" includes all the capacities made possible by the interaction of its hardware and software.

In human beings, of course, the matter is much more complicated. A complex interaction of cultural software, genetic predisposition, and environmental influences creates the entity we know as the person. The physical structure of our brain itself is altered through the acquisition of cultural skills during child-hood. If certain skills are not mastered by a point in our development, the brain will not possess the necessary equipment to produce them later. Hence the connection between the biological structures of our understanding and the

processes of social learning is closer in humans than the relation of hardware to software in any existing computer. We have evolved into creatures whose brain structure can be transformed through the processes of social learning. This is yet another sense in which it is truly in our nature to be cultural.

In order for a computer's hardware and software to interact, both must have a capacity to process information. My word-processing program allows my computer to process information, but it can do so only because it is loaded onto another program, an operating system like DOS or UNIX, that allows the computer to process software. Thus the information processing permitted by the software requires a prior information processor to employ it. Similarly, the operating system can run only because the computer has a program in firmware—read-only modules attached to the computer's architecture—that allows the computer to understand and process the commands it receives from the operating system. Finally, this firmware can operate only because the hardware of the computer allows it to process the commands of the firmware at a mechanical level. So the distinction between hardware and software in computers is not a distinction between the part of the computer that processes information and the part that does not. Rather, information processing occurs all the way down. In like fashion, we cannot say that our ability to reason and evaluate is purely a product of our cultural software. We are born with the ability to become reasoning beings. Rather, cultural software articulates, supplements, and refines our powers of reasoning and evaluation. Cultural software is the historical component of human reason, not its sole component.

The relationship between hardware and software in computers must be explained differently. In theory, my word-processing program could be hardwired into the computer. It could become part of the hardware. But in practice it is more convenient for me to be able to remove it from memory and substitute different programs, or to upgrade the program that I have. This is the great advantage of software as an information-processing device. It is changeable and adaptable; it creates the possibility of many different types of hardware/software combinations, and hence many different types of computers.²⁵

Just as computer software allows computers to harness their power, cultural software empowers human beings. The human mind is a marvelous device. But like the most powerful supercomputer, it needs methods of understanding if its power is to be tapped. Our cultural software is the result of a long process of collective accumulation and construction. It has produced elaborate tools of understanding, which, in conjunction with technology and institutions, can be tremendously empowering.26

Of course, cultural software is empowering not only in the sense of allowing us to achieve our goals. It also enables us to reflect on and describe what our goals are. Cultural software allows human beings to articulate and concretize their values, to put flesh on the bones of their innate but inchoate urge to value and evaluate. Through cultural software our brute sense of the beautiful is transformed into the many varieties of aesthetic judgment, some of which come into being and fade away at different points in history. Through cultural software the inchoate sense of good and bad is transformed into the many varieties of moral and practical judgment, and the many virtues and vices are articulated and differentiated. Thus cultural software is the great enabling device not only of human understanding but also of human evaluation. For this reason alone it is the greatest of human creations, the most powerful and important of human tools.

Historical Existence and Cultural Construction

The theory of cultural software is both a theory of ideology and a theory of historical existence. Gadamer's ontological hermeneutics argues that to exist in history is to exist in a historically generated tradition. But this answer simply raises in a new form the familiar question of what a tradition is such that people can exist within it. The theory of cultural software allows us a better way of expressing this insight. To exist in history means to be the bearer of a particular variety of cultural software that has been produced through a process of cultural evolution. Thus historical existence is not merely existence in time but existence at a time when one is constituted by a particular form of cultural software peculiar to that time. A person living in the sixteenth century has a different kind of existence from that of a person living in the twentieth, a difference that is due not merely to differences in climate and technology. Their genetic inheritance may be roughly the same, but their cultural software is quite different. And so the persons, who incorporate cultural software, are different. It is this feature of human being that distinguishes the existence of a person from that of the Grand Canyon. The Grand Canyon exists in time, but only people exist in history, because only people are constituted by an evolving, collectively created cultural software.

Like Gadamer's theory of tradition, the theory of cultural software is also a theory of understanding, or, more accurately, of the historical basis of understanding. Human understanding is understanding in history; it is made possible by tools of understanding that bear the marks of their historical development. So there is an intimate connection between historical existence and historical understanding, between living in history and understanding in history. To be human is to be constituted by a certain type of cultural software that predisposes and facilitates understanding in certain ways and not in oth-

ers—a cultural software that is the product of a certain history of conceptual bricolage and cultural evolution. This predicament is both the meaning of historical existence and the precondition of cultural understanding.

In this way, the theory of cultural software takes literally the contemporary chestnut that individuals are socially constructed. People become people only when they enter into culture, which is to say, only when culture enters into them, and becomes them, when they are programmed with and hence constituted by tools of understanding created by a culture at a certain point in history. Through existence in history, which is existence in culture, people obtain and incorporate cultural tools, and these become as much a part of them as their arms and legs.

The idea of cultural construction is often associated with cultural determinism. But the theory of cultural software suggests that this view is misleading, for cultural software empowers individuals even as it constructs and creates them. It untaps the potential power of the human mind just as an increasingly complicated and sophisticated software program allows a computer to do more. So we must understand cultural software as constitutive not only of identity but of autonomy as well. When we confuse cultural construction with cultural determinism, we misunderstand what culture does for human beings. Culture is not a law of obedience but the source of what we call freedom. Cultural software, rather than being the enemy of human autonomy, is the very condition of its possibility.

Although cultural software empowers individuals, it also creates a certain opportunity for power over individuals who are constituted by it. It does both of these things at one and the same time, and through the same mechanism. The power that cultural software makes possible is precisely the power that the tools of understanding have over the individuals who are partly constituted by them. This power arises in part from the limitations of our conceptual apparatus; this is akin to the very powerful computer that has only a checkers program. Yet a second aspect of this power is more subtle. It arises from the nature of information processing itself, and it is never fully eliminated, no matter how sophisticated the software becomes.

Processing information always requires partiality and selectivity. As Heraclitus recognized, the world is in perpetual flux; we cannot comprehend its nature in all of its infinite diversity and differentiation. Without some form of simplification, in the form of categorizations, narratives, heuristics, or norms, it is impossible to understand anything at all. Information requires simplification—taming the flux for the purpose of understanding—and so at the very moment when understanding is made possible, partiality also emerges. I often like to say that the key to information is in formation; it lies in the selection and categorization of the flux of experience into comprehensible categories,

events, and narratives. In order to understand, we must establish similarities and differences, categories and narratives, canons and heuristics. These are the basis of all information, and hence the basis of our cultural software. So our cultural software limits even as it empowers. It informs us in forming us, which is to say that it informs us in forming our selves as selves endowed with a certain form of cultural software, who see things one way and not another, who are properly "tooled up" for some tasks but not for others.

Thus cultural software has power over us because this power is rooted in the very way in which we are able to process information and articulate values. Of course, other individuals within our culture can take advantage of the partiality of our cultural software. They can gain power over us because we, like they, are constituted by the tools of understanding. The most obvious example of this phenomenon is the power of rhetoric and symbols. Rhetoric has power because understanding through rhetorical figures already forms part of our cultural software. Symbols have power because the associations that make them symbols are already part of us. So the study of rhetoric or the study of semiotics may be thought of as part of the study of cultural software, or, more properly, the study of the traces and effects of this software. It is the study of the building blocks of our understanding, and therefore the study of the forms and modes of power exercisable over that understanding. At the same time it is part of the study of reason itself, the culturally created reason that underlies our everyday thought and action.

The theory of cultural software rethinks the traditional conception of ideology in two ways. First, it sees ideological power as the power that cultural software has over the persons who are constituted by it, who are persons because of it. Instead of seeing ideology in the form of false beliefs held by subjects who preexist those beliefs, it locates the source of ideological power in the constitution of subjectivity itself. This subjectivity is not only the meaning that others assign to you but also the meaning that you assign to the world itself through the shared tools of cultural software.

Second, the theory argues that ideology, or rather what replaces it—cultural software—must be understood not only through its negative effects but also through its positive ones. Cultural software does not merely obscure; it also clarifies. It does not merely limit the imagination but empowers it as well. The theory of cultural software thus rejects a uniformly pejorative conception that views ideology as a disease or a decrepit form of human thought. In the theory of cultural software, the mechanisms of ideological thought are the mechanisms of everyday thought. In this theory, truth and falsity, deception and empowerment enter through the same door.

I have compared cultural software to a toolmaking tool. Yet to grasp the full meaning of this comparison we must understand the word tool in its broadest possible sense. The tools of understanding that I call cultural software are different from hammers and nails. Hammers and nails are made by human beings but are physically separate from them. I can pick up a hammer or put it down. I can carry it with me or leave it at home. Not so with the tools of understanding. The tools of understanding work by becoming part of my apparatus of understanding, which is to say they work by becoming part of me. Cultural software is not just something that we use to understand and evaluate the world; it is also part of us. Indeed, human beings do not become persons until they enter into culture and become imbued with some form of cultural software. To exist as a person is to exist as a person who has cultural software, who is, in part, her cultural software.

There is an old anthropological view of humanity as homo faber, Man the Toolmaker. It argues that human beings are distinctive because they fashion tools for their own use. Its paradigmatic conception of human activity is technical, and its paradigmatic vision of human reasoning is instrumental. This conception is too narrow to provide a just account of the human condition.² It focuses on the acquisition of technical skills rather than on the social skills that are central to our existence as cultural beings. Yet the limitations of the homo faber thesis stem not so much from its view that human beings fashion tools but from its limited conception of what tools are. And for many people the idea of tools does not go much beyond this narrow conception. Because the metaphor of the tool is central to my argument, it is important to attack this conception and show its limitations.

What Is a Cultural Tool?

The narrow conception of tools that I find objectionable rests upon four implicit assumptions which at first seem almost definitional: First, a tool is a material object that is physically separable from the person who wields it. It can be taken up or discarded at will. Second, a tool is designed for a particular purpose and has a specific function. Third, the tool's use is purely instrumental—to further a preexisting end specified by its user. A tool is, quite literally, an instrument, and hence a pure adjunct of instrumental rationality. Fourth, using a tool involves a purely technical skill.

At first, these four assumptions seem perfectly natural and even necessary, especially if our paradigmatic examples of tools are hammers, knives, and clay pots. Yet hammers, knives, and pots were not the only tools that humankind developed. People also developed language and other social skills. Once we recognize language as the quintessential cultural tool, each of the previous four assumptions must be jettisoned. First, our linguistic abilities cannot be separated from us. They are part of who we are. Second, these abilities are not designed for any one purpose. We use language for any number of purposes. Indeed, we use language to formulate new purposes for which language will subsequently be used. Third, our use of language is not purely instrumental; we use language to express our values and interact with others. Finally, our use of language is not a purely technical skill; it is the social skill par excellence.

The tools of understanding that I am primarily interested in fall into this latter category. They are internal to and inseparable from human existence. They prominently include social as well as technical skills. They are not necessarily designed for a single purpose but have multiple purposes and are often the source of new purposes. They are not simply means to an end but the means of developing and articulating our ends.

In fact, once we recognize language, and not the clay pot, as the paradigmatic cultural tool, we begin to realize that even material objects like hammers and pots do not fit the limited conception of tools that I have associated with homo faber. Material tools do not have to be designed for or serve a single purpose, for example. Lévi-Strauss pointed out that much human activity is like the work of a bricoleur, or odd-job man, who takes whatever implements lie to hand and uses them for new and unintended purposes.

Second, even if material objects are separable from our physical self, they can be important and constitutive parts of our social self. Hegel pointed out that our social identity is strongly linked to our possessions, the ways we use them, and the opportunities they make available to us. As King Lear discovered, to abandon one's property is simultaneously to radically change one's social relationships. The person who surrenders her glasses, her telephone, her car,

and her computer changes not only her instrumental abilities but also her social life. The destitute and the homeless may lack certain instrumentalities, but the losses they feel are not purely instrumental ones.

Third, people use their intelligence to create tools, but these tools simultaneously endow their users with new kinds of intelligence, because they allow users to experience and interact with the world in new ways.³ The tool changes the horizons of our imagination, opens up new perspectives, and makes new kinds of considerations possible. The tool does not simply change our ability to experience the world; it also changes the world we experience.

Tools change our imaginative horizons because the world now appears as a possible object of manipulation in new ways. Yet tools also change our perspectives because they help create a world centered around their use and the products of their use. Tools like computers, looms, and printing presses also give rise to new concepts, new skills, and new concerns. They not only make it possible to analyze the world in new ways, but they also become objects in the world around which new social skills, social institutions, and artistic practices can be organized. Thus the development of the hammer makes possible the development of skills involving a hammer, like forging iron. The invention of the automobile gives rise to the possibility of magazines, clubs, and institutions organized around the manufacture, racing, display, and general admiration of cars. The development of musical instruments gives rise to the possibility of organizations like symphony orchestras, professions like conducting, and eventually, to the rise of a music industry, schools of performance, and the practice of music criticism. Computers make possible computerassisted graphics and design, electronic churches and chatlines, and books like this one, whose central metaphor compares cultural know-how to computer software.

Finally, tools are not always mere adjuncts of instrumental rationality. They are used in many different ways. In particular, I want to distinguish three different uses of cultural tools. The first is to get about the world, to understand and make use of it. The second is to interact with other people, and the third is to express and articulate human values.

In practice these purposes surely overlap. I distinguish them analytically because I want to contest the natural association of tools and toolmaking with the first purpose, and in particular, with instrumental rationality. The slide from instrument to instrumental is easy to make; too often we think of tools solely as a way of exploring and mastering the natural world. This mastery can be either material or intellectual—it may occur either through controlling and shaping nature or through understanding it. Nevertheless, if this were the only point of toolmaking, the conception of culture as a set of tools, and the conception of humankind as a toolmaker and tool user, would be significantly

impoverished. Human culture would be thoroughly instrumental, nothing more than a means by which human beings master their environment. Such a conception of culture would fail to recognize the existence of other human beings; or, in the alternative, it would view them as just another set of objects to be controlled, governed, studied, and mastered. Similarly, the concept of reason developed through culture would be reduced to instrumental rationality. People would not be able to reason about values or ends but only about means.

We may justly criticize a culture to the extent to which it is primarily or excessively concerned with instrumental rationality at the expense of other forms of reason. Such a view forms the basis of Horkheimer and Adorno's critique of the Enlightenment.⁴ Nevertheless, we should not confuse this criticism with a criticism of the metaphor of toolmaking or with the conception of human beings as toolmakers. That criticism is valid only if toolmaking really has no other purpose than the mastery of objects. The tendency to think that this is so may itself be a symptom of living in a culture that has placed too high an emphasis on instrumental concerns and instrumental rationality.

A second and quite different purpose for making and using cultural tools specifically concerns our relationships with others. Culture and cultural tools enable us to treat other individuals as persons and to negotiate (and struggle over) shared meanings with them. Language is perhaps the best example of this sort of tool. This use of cultural tools presupposes the existence of other minds that are recognized as others, rather than merely as objects of control. The ideas of negotiation and struggle require other thinking beings with whom to cooperate or contend. This recognition is as true of cooperative ventures as it is of competition and even war. We have a relationship with an enemy, because an enemy is another person against whom one struggles, rather than merely an object or a force to be mastered or dominated.⁵ Thus all forms of human conflict, including war, have both instrumental and intersubjective features.

An important feature of cultural toolmaking, then, is to share with, negotiate with, and struggle with others as others. We use cultural tools to communicate with others, play with others, care for others, work with others, and fight with others. Much of our technology involves devices for interpersonal interaction, cooperation, and struggle; examples include telephones, baseball gloves, bedpans, mascara, and machine guns. The very word discourse that has emerged as a substitute for the concept of ideology implies an intersubjective connection between ourselves and others. Discourse comes from Latin words meaning to run back and forth. It implies a bidirectional movement between parties rather than a unidirectional control of an object by an intelligence. In the same way, conversation comes from a Latin word meaning to turn around also implying a notion of reciprocity.

There is some irony, therefore, in Michel Foucault's adoption of discours

to explain cultural power. He writes, for example, of the emergence of "a discourse in which the sexual conduct of the population was taken both as an object of analysis and as a target of intervention."6 Here discourse becomes a technology of governance over a population, a way of normalizing, shaping, and controlling people's behavior. Often Foucault even speaks of individuals in terms of bodies that must be directed and controlled through discourse as well as technology. Thus, for Foucault, discours often has a strongly instrumental flavor, although it is by no means clear in his writings who, if anyone, is wielding the instrument.

The normative bite of Foucault's analysis of discourse stems precisely from the way in which discourse controls people and their bodies, objectifies them, and denies their freedom—in other words, precisely from the ways in which discourse is false to the idea of communicative reciprocity. On the other hand, sometimes Foucault seems to speak as if it were appropriate to analyze selves as merely the products of discourse, and the objects of control by discourse (as well as technology). This subverts the normative uptake of his analysis. I argue instead that to the extent that culture merely makes people into objects of control, it can justly be criticized for transforming all cultural tools into instrumental tools and all persons into objects of control or governance. Indeed, as I shall argue more fully in Chapter 12, if culture is not understood in terms of its intersubjective aspects, it is difficult to account for struggle and resistance, because individuals become simply the intersections of larger forces of discourse and technologies of bodily control as opposed to situated agents who employ cultural tools in their struggles.

Nevertheless, it is important to recognize that the intersubjective aspect of culture, like the instrumental, is by no means uniformly benign. If language is a tool used to cooperate with others, it is also one of many tools that human beings use to struggle with and dominate others. Our distinction between the instrumental and the intersubjective aspects of culture is not a distinction between the harmful and the beneficial uses of culture; within each perspective the tools of culture can be employed for good or for ill.

Cultural Software and the Articulation of Values

A third use of culture and the tools of culture is the articulation and expression of human values. Human beings have values, and these values are one of the most important features of human life. Or more correctly, human beings value, for we should think of value primarily as a verb, not a noun. Values are not so much what people have as what they do and feel. Human beings possess an inexhaustible drive to evaluate, to pronounce what is good and bad, beautiful and ugly, advantageous and disadvantageous. Without culture, human values

are inchoate and indeterminate; through culture they become differentiated, articulated, and refined.

Consider the example of music. Before culture there are no electric guitars, violins, or orchestras. There is no art of orchestration, no sonata-allegro form, no idea of jazz or the blues. There is only the human delight in producing and listening to interesting and beautiful sounds. Throughout human history people develop different ways of making and organizing sounds, which they test against their developing sense of beauty and interest. Their sense of the beautiful and the interesting in turn is developed through exposure to and use of the cultural tools available to them within their culture. So a person who lives in Austria in Mozart's time has a certain sense of taste about what is delightful and interesting in music. Her values are articulated in a certain way, although in another culture and another time they might be articulated differently.

The word articulate comes from the Latin articulus, meaning organized in joints or joined. It carries the double meaning of dividing something up into distinct parts (hence a person is articulate who can make distinct sounds) and joining parts together (for example, the articles of a larger document, like the Articles of Confederation, an early form of government for the United States). Thus articulation involves both distinction and construction; it consists in both the refinement of old values and the creation of new values from old ones.

Both refinement and construction are involved in the cultural articulation of musical tastes. People develop their tastes by becoming able to make distinctions between different pieces of music or different performances of music. If we have never heard jazz before, we may have only a vague idea of what we like and dislike. We may not be able to distinguish Coleman Hawkins from John Coltrane, or Thelonious Monk from McCoy Tyner; in addition, our sense of what is enjoyable and what is tedious may be limited or inchoate. After we hear more jazz, two things may happen. First, we become accustomed to jazz and we may come to enjoy compositions we would not have previously enjoyed. Second, we become increasingly able to distinguish between different performances of jazz and can express our judgments with greater distinction and refinement. The notion of refinement involves both the alteration of our tastes through exposure and familiarity (like the development of a taste for wine) and the ability to recognize distinctions in what previously seemed an undifferentiated whole. Hence we say that a person's taste is refined precisely when she is able to make distinctions, separating out the better from the worse.

Yet cultural articulation also involves construction. Culture does not merely enable us to make increasingly finer distinctions; it also enables us to create new possibilities for musical enjoyment and musical evaluation by creating new types of instruments, new forms of musical expression, and new musical compositions. These cultural constructions are passed on and modified from generation to generation. They become part of our developing sense of musical taste and enable us in turn to make new evaluative distinctions, distinctions that were not previously possible because they partly presuppose cultural constructions that had not yet come into being. In this way culture continually creates new tools for musical evaluation and expression.

Like our aesthetic sense, our ethical sense is also articulated through culture, though the terms of this articulation are by no means identical in all respects. Broadly construed, our ethical sense concerns how we should live our lives, as well as the evaluation of what is praiseworthy and what is deserving of scorn. An example of this articulation is the historical proliferation of virtues and vices. By this I do not mean that people become more virtuous or wicked by living in culture. I mean that through culture we come to divide up simpler notions like good and bad into a kaleidoscope of varieties and modes of virtue and vice. Through culture we come to understand many different varieties of good human character and activity, including wisdom, mercy, friendliness, loyalty, courage, and justice. At the same time, we come to recognize and distinguish many different bad aspects, including snobbishness, hostility, indifference, sloppiness, conceitedness, avarice, and aggressiveness.

Just as there is a proliferation of evaluations of human character and moral activity, there is a proliferation of good and bad human qualities that are not virtues or vices (for example, being impoverished or honored), good and bad features of human institutions, and good or bad features of inanimate objects. Indeed, a large part of human language is devoted to evaluative concepts that articulate, refine, and subdivide the inchoate ideas of the good and the bad. Through culture people articulate their evaluative sense into different conceptions like good and evil, pious and impious, advantageous and disadvantageous, fortunate and unfortunate, healthy and unhealthy, beautiful and ugly, sublime and mundane, noble and base. These categories, in turn, are further divided, constructed, refined, reconceptualized, and replaced. Hence, from more simple, indeterminate, and inchoate human values culture constructs a complex, rich, and detailed language of human evaluation.

Like the articulation of musical tastes, the articulation of human virtues and vices occurs within a cultural and historical context and hence develops and changes with that context. Compare, for example, the treatment of artificiality and naturalness between the eighteenth and nineteenth centuries. In the Age of Reason, artificiality is a virtue; it illustrates the application of human intelligence to a situation. In the Romantic age, artificiality is viewed as a vice, and naturalness, lack of pretension, and spontaneity (all identified in one way or another with authenticity) are exalted as virtues. In this way each age and culture articulates various features of the human condition that are worthy of praise or blame.

By this example I do not mean to suggest that every act or every trait that is virtuous in one era or place will necessarily be viewed as harmful or wicked in another, or that there is no common ground between cultures and times. Rather, differences in associations may occur because a culture places a different emphasis on traits and behaviors, in part because of the history and cultural context bequeathed to them.

Moreover, because evaluative terms are articulated through cultural development, some evaluative concepts come into being or gain greater importance at certain points in history, while others become less important, are largely forgotten, or even fade away. Thus the idea of being neurotic, obsessive, or compulsive is largely a concern of our own time, while speaking of people as being temperate or intemperate is a less common way of thinking about and evaluating them, though it was quite important in an earlier age.⁷

As in the case of our aesthetic sense, the articulation of our ethical sense involves not only distinction but also construction. Some virtues and vices are internal to particular practices that arise only at particular points in human history. Thus one cannot have the vice of being a promise breaker until there is a practice of promising; one cannot have the virtue of being creditworthy until there are financial institutions that bestow credit.

Culture also enables human beings to express their values through construction and exemplification in concrete institutions, practices, and behaviors. Take, for example, the value of justice. This value is inchoate and indeterminate. In order to give meaning to it in our lives, we must exemplify it in institutions, rules, or a system of law. Not only is this construction necessary to achieve justice, it also provides us with necessary tools and examples for further reasoning about what is just and unjust. The importance of concrete exemplification to theoretical development is likely to be overlooked. In the Republic and the Laws, Plato tried to define justice. But he discovered that in order to explain his vision, he had to construct an imaginary state, complete with institutions, social classes, occupations, offices, and regulations. The same is true in the world outside philosophical speculation. We concretize our indeterminate value of justice by creating human institutions and practices that attempt to enforce it and exemplify it, even (and especially) if we recognize that all of these institutions are imperfectly just. Of course, because justice is an indeterminate standard, there is no necessary way to exemplify it. The value of justice does not tell us, for example, whether a democratic legislature should have one, two, or three houses. Hence the institutions that people construct to exemplify justice may be different in different eras and different lands.8

It follows from the same line of reasoning that human beings can also generate ever new examples of injustice and oppression through their cultural constructions. In different times and places, human beings find new ways to work evils on their fellow creatures, and to create monuments to brutality and repulsiveness. Thus, when I say that culture allows us to refine and articulate our values, I do not mean that culture necessarily makes us better people or leads unequivocally to what is good. Moreover, when I say that people use culture to exemplify their values through constructing practices and institutions, I do not mean that whatever standards a culture produces define what is good and bad, or that it is not possible to criticize a culture for producing wicked practices or unjust institutions. I endorse neither a claim of progressive betterment through culture nor a claim of pervasive moral relativism. Rather, I am arguing that people use culture as a kind of tool to express their values and to put flesh on the bones of their inchoate urges by constructing concrete examples of what they value. Of course, like all tools, these cultural tools can be used well or ill, skillfully or awkwardly, and what they generate can be well or ill produced. If we think that culture can develop or refine our tastes, we must concede that it may also debauch or coarsen them.

Bricolage and the Creation of Cultural Software

Cultural software consists of collectively created tools that constitute us as persons and that we use to make new tools for understanding the world around us, interacting with others, and expressing our values. Yet one cannot make something out of nothing. The tools that we create must be constructed out of those we already possess. We must make all our new cultural tools out of our old ones: this is as true of our cultural software as it is of our technology and our institutions.

The history of thought is the history of the cumulative marshaling of existing capacities to form new ones, the use of older cultural software to create newer "idea-programs." This process of cumulative construction of new conceptual tools out of old ones resembles Claude Lévi-Strauss's notion of bricolage. Lévi-Strauss argued that human thought operates like the bricoleur, or odd-job person, who fixes a leak or other problem with whatever tools lie to hand.9 My use of bricolage differs in an important way. An odd-job person repairs leaky faucets and roofs that remain with the customer after the repairman moves on; the bricoleur repairs them with tools that are distinct from the roof and the leaky faucet. In my conception of bricolage, what the bricoleur creates in her jerry-built fashion she keeps with her to use in the next job, and the next. The products of earlier bricolage become the new forms and methods of later bricolage.

The claim that culture, and in particular cultural software, is the result of

bricolage entails four basic notions. Cultural bricolage (1) is cumulative, (2) involves unintended uses, (3) is economical or recursive, and (4) has unintended consequences. Let us consider each of these features in turn.

First, cultural bricolage is cumulative. The tools of understanding that one can create at a particular time depend largely on the available materials that lie to hand. The complexity and performance of a tool are necessarily limited by the nature of the tools available to construct it. Consider the examples of a spacecraft and a mutual fund. A spacecraft requires the development of sophisticated forms of metallurgy, the production of powerful chemical fuels, and the construction of elaborate electronic systems for computation and communication. The construction of the Apollo spacecraft that traveled to the moon in 1969 was not just the result of a huge investment in engineering skill during the 1960s but also depended on centuries of innovation that made the final stages of technological development possible.

A mutual fund is a device for lowering the risk of investing in bonds, securities, and other financial instruments. Yet in order to create such a fund, elaborate financial institutions already have to be in place, which, in turn, could only be developed after the creation of previous financial institutions. The capitalist "free market" involves much more than simply allowing people to buy and sell commodities, as the reformers in Eastern Europe discovered soon after the fall of communism in 1989. It requires elaborate institutional mechanisms for amassing and distributing capital, spreading financial risk, and developing new instruments for capital investment, institutions that required centuries of experimentation and development.

Just as human beings engage in cumulative development of technology and institutions, they engage in cumulative development of their cultural software. Through this development increasingly rich and varied modes of thought become available to human beings. Conversely, the ability of human beings to articulate new ways of thinking depends upon the cultural inheritance bequeathed to them and upon the tools of understanding available at a particular point in their history. This development is often nonlinear and unpredictable; developments in one area of culture may hold the key to changes in a completely unrelated area. Cultural development is the unanticipated use of the unexpected, passing under the name of rational progress.

A second feature of cultural bricolage results from the first. Because the bricoleur uses whatever tools lie to hand, she does not necessarily employ tools for their original purpose. She may employ a screwdriver as a makeshift hammer, or use a bucket as a doorstop. Thus, central to the concept of bricolage is the possibility of unintended use. By this I do not mean that tools are not used intentionally by agents but rather that they are put to purposes for which they were not originally intended. Like other tools, human institutions may

also be adapted to new purposes originally unintended. The organization of the family, for example, can be employed by analogy to organize a religious movement, and the features of religious organizations may be adapted to political or social organization. Democratic political concepts, in turn, can be grafted onto the structures of religious organizations. The mechanism of a market can be applied in countless ways. In this fashion human institutions solve problems of organization, reproduction, and stabilization by adopting and adapting features of other social structures that their members are familiar with. In this way new forms of human sociability are constructed out of older ones.

Cultural software is also the product of unintended usage of previous concepts. People use familiar concepts in order to describe the world and construct new concepts. Wittgenstein explained the nature of language and thought, for example, by comparing it to a game. Once this is done the notion of a "language game" takes on a meaning of its own and can be used metaphorically or analogically by later thinkers.¹⁰ This book is itself an exercise in conceptual bricolage, for it borrows from several disciplines in order to construct its argument.

Many examples of previous conceptual bricolage can be discovered through the study of etymology. For example, I noted earlier that the word articulate comes from the Latin articulus, meaning joined or jointed. If we pursue this etymology further, we discover that articulus and ars (art) have a common ancestor. The concept of art may originally have developed from the idea of joining or assembling something. Both of these words, in turn, probably share a common ancestry with the English word arm. Indeed, it is possible that articulus and ars are metaphorical extensions of an ancient word for arm. The arm is the most obvious example of something that is joined to the body and that itself contains joints. In this way words that are used for one purpose are extended to serve new purposes. Thus from ars we get artist, artisan, artful, and artless, and from articulus we get not only articulate (to divide sounds, hence having the ability to do so, hence eloquent) and article (a division of words, hence both a grammatical form and a type of writing) but also arthritis (a condition of the joints). Human language often develops new concepts by metaphorical, metonymic, or analogical extensions of older concepts; these concepts, in turn, may be employed to develop still newer ones that bear only the remotest relationship to their ancestors. 11 In this manner language moves from the concept of an arm to the idea of eloquence.

A third feature of cultural bricolage is its economy: a relatively small number of tools are used in many different situations to do a comparatively large number of jobs. The social theorist Pierre Bourdieu has used the expression "economy of logic" to describe this phenomenon. 12 Bourdieu points out that cultures can use a relatively small set of conceptual oppositions repeatedly to

generate increasingly complex sets of cultural meanings. Gender is a good example of this process. In many cultures, including our own, the concepts of male and female are assigned not only to human bodies but to many types of concepts and objects. First, there are familiar associations regarding men and women. In our culture, for example, there are stereotypes involving colors (blue and pink), tools (hammers and brooms), food (steaks and salads), and literature (adventure and romance). Second, in various cultures, objects and concepts may be called male or female—for example, Father Time and Mother Nature. Inanimate objects may be labeled male or female because of analogical comparisons to male or female anatomy (male or female phone jacks) or to stereotypical "male" or "female" traits under patriarchal ideology (the moon as the lesser light, the reflection of the sun). Third, and most important, gender categories may be assigned simply for purposes of conceptual bookkeeping and division. Examples are languages that divide all nouns into male, female, or neutral genders. Grammatical assignments of gender are a good example of the cumulative nature of bricolage, for often these assignments have only a very indirect relationship to historical male and female stereotypes. More often they have none at all; it is simply necessary as a matter of linguistic convention that everything be assigned a gender. In this way the grammatical use of the term gender comes to have an entirely new meaning.

The cumulative use of gender categories in diverse situations for a vast variety of purposes thus results in ubiquitous male and female encodings in language, thought, and practice.¹³ Yet this is not the result of deliberate design. It is rather the result of the repetitive use of simple tools to fashion newer tools, all of which bear the marks of the previous tools used to make them, and each of which transmits this marking to the tools that it in turn is used to make. Conceptual bricolage is a repetitive and recursive semiosis or meaning making. It is repetitive because it is used in many different contexts and for many different purposes; it is recursive because it is applied to results of previous conceptual bricolage. Because conceptual bricolage is repetitive and recursive, there is an important connection between the economy of bricolage and the phenomenon of unintended usage. The conceptual bricoleur uses concepts, distinctions, and frameworks repeatedly in new situations; this virtually guarantees that these concepts, distinctions, and frameworks will be employed in contexts and for purposes for which they were not originally designed.

A fourth feature of bricolage results from the previous three features. The bricoleur's economical and cumulative use of tools in unintended ways can and often does lead to unexpected and unintended consequences both for good and for ill. This is perhaps the single most important idea in the philosophy of culture—the unexpected consequences of human thought and action. It rightly plays a central role in many philosophies of history and culture. Vico's idea of Providence, Kant's conception of Nature in the history of human development, Hegel's notion of the "cunning of history," and Marx's claim that men make history, but not as they intend it, all exemplify this insight. The concept of cultural bricolage is yet another way of approaching this difficult but fundamental idea in the philosophy of culture.

The unintended effects of cultural bricolage can be both positive and negative. In fact, they tend to be both positive and negative at the same time. They have simultaneous unexpected and unplanned benefits and disadvantages. We can better understand this phenomenon through an analogy to another form of development, biological evolution.

Cultural Bricolage as a Form of Evolution

The idea that cultures evolve is, if anything, older than the idea that species evolve. 14 The rise of Darwinian theory made it natural to see possible analogies between cultural and biological evolution, and various attempts have been made to do so from Darwin's time to the present day. Such analogies, however, can be misleading unless one notes the grounds of difference as well as similarity. We might begin by distinguishing between Darwinian and Lamarckian theories of evolution. Lamarckian evolution argues that organisms survive by adapting to their environment and passing on their acquired characteristics to their offspring; Darwinian evolution argues that organisms that have relatively adaptive characteristics are more likely to survive and pass their genes on to their offspring. A Lamarckian would argue that giraffes' necks became long because giraffes kept stretching them in order to reach high leaves on trees; a Darwinian would argue that giraffes with long necks were better able to survive in times of limited food supplies. In Lamarckian evolution, variation occurs as a response to the environment; in Darwinian evolution, variation is random and the environment weeds out the comparatively maladapted.¹⁵

Biological evolution appears to be Darwinian, employing natural selection as its central mechanism.¹⁶ Cultural evolution, however, seems to be both Darwinian and Lamarckian.¹⁷ Cultural evolution does not proceed merely through recalcitrant experience choosing among various cultures and acculturated individuals. Members of a culture can to some degree self-consciously understand the problems facing them, change aspects of their culture to face new challenges, and pass these changes in cultural software, institutions, and technology to succeeding generations.

This description of cultural evolution is not, strictly speaking, Lamarckian. Unlike the giraffes in Lamarck's theory, human beings do not pass on their acquired cultural innovations through their genetic materials; they pass them on though social learning. To say that cultural evolution is Lamarckian, we must make two important additional assumptions. First, we must regard human beings as combinations of both their genes *and* their cultural know-how. Second, we must regard social learning as a form of nongenetic inheritance. If we make these two assumptions, cultural evolution becomes the evolution of cultural know-how transmitted through culture-carrying creatures. These creatures evolve because the cultural component of their being varies and develops over time.

Because cultural evolution involves cumulative social learning, it can proceed much more quickly than biological evolution. Biological evolution has taken hundreds of millions of years to produce intelligent life on this planet, while cultural evolution has taken thousands of years to bring human abilities to their present levels.

The conception of culture as a set of toolmaking tools is Lamarckian in the sense that cultural know-how can become part of people and can be passed along to succeeding generations through social learning. Nevertheless, an evolutionary theory of culture, whether Lamarckian or Darwinian, faces a significant problem. Both types of theories assume that evolution is produced by differential rates of survival for entities in a given environment. If cultural evolution is a kind of human evolutionary process, individuals or groups of individuals that adopt certain innovations would tend to have greater chances of survival in a particular environment than those that do not. 19 But it is difficult to explain many aspects of cultural change in terms of their increased survival value for particular human beings or for the human species generally. Although some forms of cultural development do benefit human beings by increasing their chances for survival (for example, the development of medical science), many others do not. There is no reason to think that the proliferation of different guitar and violin designs, for example, or the successive fashions and styles of popular music or dress significantly assist the survival of the human species.

I wish to make two basic claims about the relation between conceptual bricolage and theories of evolution. First, for the reasons just stated, the process of conceptual bricolage is not, like biological evolution, a form of natural selection of human traits. The content of human cultural software is not driven by differential rates of human survival in the same way as human genes. Nevertheless, I shall argue in this chapter that there are many important analogies between the historical process of conceptual bricolage by human beings and the biological evolution of organisms, and these comparisons shed considerable light on the nature of cultural software. In drawing these analogies I am working in the opposite direction from a famous paper on evolutionary biology by Stephen Jay Gould and Richard C. Lewontin.²⁰ They used examples of cultural bricolage to elucidate the mechanisms of biological evolution. Gould and Le-

wontin harbored no illusions that the processes of cultural and biological development were identical; but they also understood that these differences did not foreclose important points of comparison.

My second major claim about culture and evolutionary theory is the subject of Chapters 3 and 4. Although the growth of cultural software is not a form of natural selection of human beings, it is a form of evolution. There is a genuine Darwinian process involved in its development. However, this process is not primarily concerned with human survival, and the unit of selection is not human genes, human beings, or groups of human beings. Instead, the unit of selection is cultural software itself, and the "environment" in which it competes, survives, and propagates is the human mind.

Analogies Between Cultural Bricolage and Biological Evolution

As I have argued, one of the most important features of conceptual bricolage is its connection to the unintended consequences of human thought and action. We can draw several analogies between the evolutionary development of organisms and the many important and powerful features of human culture that are not the product of conscious planning.

First, let us consider the question of unexpected advantages. We might begin with a distinction between designed and designoid objects.²¹ Designed objects are the result of conscious shaping and planning by some intelligence; designoid objects are objects that appear to be designed but are actually the unintended or nonintentional consequences of causal forces. An example of a designoid object would be the symmetrical pattern of a crystal, or the distribution of iron filings when they come into contact with a magnetic field. Darwinian evolution assumes that all living organisms are designoid. The human eye operates and functions as if it were designed; indeed, it operates better than any device created by human engineering. However, Darwinian evolution holds that this is the result of natural selection and other evolutionary forces as opposed to conscious planning.²²

Although much of human culture is the product of conscious design, many other aspects of culture can be described as designoid. For example, consider a market. A market sets prices for commodities, takes orders for production, and distributes goods and services. All of these tasks might be deliberately planned and performed under a command economy. But a market achieves them without the conscious design of any individual or set of individuals. Thus a market, although making use of the plans and intentions of individual agents, produces a system of production and distribution that is not designed but designoid. Indeed, advocates of laissez-faire argue that a designoid market performs better than a consciously designed command economy.

Just as environmental features work like a refiner's fire to produce magnificent structures that would be the envy of any designer, the cumulative development of culture through bricolage can have unexpected benefits. Nevertheless, this presents an overly one-sided view of biological evolution. Not all features of organisms, even those of very successful organisms, are perfectly arranged. Nor is every feature of an organism maximally adaptive for its present environment. Indeed, if this were the case, it would cast serious doubt on the Darwinian theory of evolution, or, for that matter, any sort of evolutionary explanation. To the contrary, the best evidence of evolutionary forces is that certain features of organisms are imperfect and poorly crafted.²³ These imperfections provide us with another analogy to the process of cultural bricolage.

The evolutionary theorist Stephen Jay Gould offers the giant panda as an example of evolutionary bricolage. The panda has a bone extending from its wrist that acts as a primitive and clumsy opposable thumb; this thumb enables it to eat its staple diet of bamboo. Compared with the human thumb, the panda's is awkward and poorly designed. This awkwardness, though, is evidence of evolutionary development. The herbivorous panda evolved from carnivores that used their digits for clawing, not grasping. The previous evolutionary development of carnivores foreclosed the development of a digit into a thumb; instead, a "thumb" developed as an extension of the wrist bone.²⁴

The panda's thumb exemplifies several important features of evolutionary explanation. Evolution makes do with the materials it is handed. What it has to work with depends upon the evolutionary problems of the past and previous responses to them. As a result, an evolutionary mechanism cannot always employ the solution that would be best if an organism were to be designed from scratch. Rather, the very awkwardness or imperfection of a solution to environmental pressures is evidence of the historical nature of evolutionary production. Organisms carry their history with them, so to speak, and this history shapes and directs the possible accommodations that an organism can make for the future.²⁵

In the same way, cultural bricolage must construct new tools of understanding out of previous ones. This shapes and constrains the way new tools are constructed. Tools are crafted to deal with the problems of a particular time and particular circumstances. They may work well enough for one set of circumstances but less well in new situations, and they may have unexpected benefits in still others. Hence there are two different sides to cultural bricolage. First, as contexts change, older tools turn out to have unexpected side effects and even deficiencies. This is analogous to the carnivore's claw that becomes less useful when placed in the new context of a herbivore's diet. Second, people may modify or adapt existing tools for purposes quite different from their original use in order to deal with new situations. This is analogous to the devel-

opment of a wrist bone to create an opposable thumb. These modified tools, in turn, will have unexpected side effects when they are placed into new contexts.

Central to the concept of cultural bricolage, then, are the simultaneous degrees of constraint and freedom produced by historical development. Although historical development always forecloses certain possibilities (in the same way that the panda's evolution from a carnivore foreclosed the development of a thumb from a digit), at the same time it also creates new possibilities for innovation (like the development of a thumb from a wrist bone). In the same way, features of existing technology, institutions, and cultural software are always potentially available for adaptation to new and unintended purposes.

Yet this freedom comes at a price. Cultural tools produced by bricolage never work perfectly: when they do work it is usually only well enough for the purpose at hand. This is as true for the predecessors of present-day cultural tools as it will be for the future products of culture. Thus the development of culture is not simply a falling away from a previous time when cultural tools were perfectly adapted to the world. There is never a time when the products of cultural bricolage lack a certain jerry-built character, when they do not have unexpected side effects or the potential for such side effects. The history of the development of culture is always the history of muddling through, with so many unexpected turns and twists along the way that "the unexpected" threatens to become the rule rather than the exception.

Existence in history produces the marks of history. Organisms produced by evolution display the remnants of previous development, which may have little relevance to the environmental problems they currently face. This accounts for so-called vestigial organs like the human little toe or appendix; the continued presence of these organs is evidence of previous evolutionary development.²⁶ Precisely because organisms do evolve, one cannot infer that because an organism currently has a certain feature, that feature is currently adaptive. Rather, one can infer only that the feature was at some point relatively adaptive or relatively advantageous (or was genetically linked to such a feature) given the particular environment in which it developed, and that it does not create such a great hindrance to the organism in its current environment as to have been eliminated through natural selection.

Indeed, if organisms are truly the products of historical development, current utility is neither a necessary nor a sufficient condition for an organism's possessing a given feature. Some existing features of an organism may have developed for another purpose but have turned out later to bring unexpected advantages to future generations. These instances of evolutionary bricolage are called extapations.²⁷ Gould points out, for example, that "feathers work beautifully in flight, but the ancestors of birds must have developed them for another purpose—probably for thermoregulation—since a few feathers on the arm of a small running reptile will not induce takeoff."²⁸ Indeed, a particular feature may have served a series of different purposes, each leaving a mark on its development. As a result, it becomes difficult to see how the feature could have evolved directly to serve the function it now serves.²⁹

Like organisms, cultural tools bear the marks of their own history—the seams, inconsistencies, and imperfections that are evidence of bricolage. 30 Like certain features of organisms, the tools of understanding may lie dormant for many years until they become useful for a new and unexpected purpose. I have previously noted that the theory of cultural software is a theory of existence in history; to exist in history is to consist in part of the cultural software developed at a particular time. Here is another way of understanding the consequences of this claim: cultures, and the people composed of cultural software who live within them, also display the remnants of previous development, which are the result of problems faced previously in the past but which may bear less relevance today. We see this in the etymology of words and in traditional practices and concepts that seem to have outlived their original use. Nevertheless, aspects of culture can always be turned to new purposes in new situations. Features of culture developed for other purposes can turn out to have unexpected uses. New ideas can be developed out of older ones by metaphoric or metonymic extension; new institutional matrices can be created out of old ones employed in different situations and times. Through this process cultural tools come to bear the marks of the previous purposes for which they have been employed. Cultural bricolage wastes little, uses much, and multiplies its imperfect improvisations on imperfect cultural tools endlessly.

I have identified the imperfections of bricolage with the application or modification of older tools in new and unexpected contexts. But the problem is implicit in the very concept of a tool. No tool is perfectly adapted for all situations and all tasks. All tools, even well-designed ones, involve trade-offs that are integral to their design and performance.³¹ An automobile is very good at traveling on land for the same reasons that it is wholly inadequate for traveling on water. Although improvements in technology can produce an amphibious vehicle that travels equally well on land and water, features of its design will make it inadequate for still other purposes, for example, travel by air, production of food, or mathematical computation. Thus the usefulness of tools is always tied to the context of their use.

Like biological evolution, cultural bricolage makes do with the available tools of understanding to create new ones. It has features of both the designed and the designoid. It is the work of human intelligence but has unexpected consequences; to paraphrase Marx, people make culture, but not as they intend. For this reason, human culture does not produce technologies, institutional

frameworks, or ways of thinking about a problem that would necessarily be best if one could design them from scratch. Indeed, the cumulative nature of historical development precludes this, for the tools necessary to engage in such a design do not exist at every point in time, and whatever human beings can create at any point in time is constrained by previous technology, institutions, and cultural software already in place.

We may offer one final comparison between cultural bricolage and biological evolution. Both are nonteleological theories of change, or are agnostic about teleology. Species evolve in response to the conditions they face, the features they currently possess, and the stock of genetically transmissible variations available at a given time. There is more than one way for species to meet a given environmental problem, and different species have solved similar problems in different ways. Some paths eventually lead to a dead end because of unforeseeable changes in environment, an insufficiently flexible set of morphological features inherited from the past, or an unluckily limited set of variations available at the time when a crisis of survival arises. Moreover, those changes that do occur need not be the most adaptive or even the best from some other normative standpoint. Rather, a change must, either by itself or in combination with other traits with which it is linked, be sufficient to guarantee the survival of the species in its present competition with other species and in the context of the local environment. Biological development is thus cumulative, but it is not necessarily a cumulative improvement. Darwinian theory is a theory of evolution away from previous conditions rather than toward a particular goal. This picture does not assume that there is some goal of increasing perfection toward which species strive; it does assume that the course of this evolution is checked by recalcitrant experience in the world. Thus this sort of evolution responds to the environment without being teleological.

By analogy we might argue that the development of cultural software is also nonteleological; cultural change does not occur as the result of a conscious plan by a unitary intelligence or the working out of an inherent natural tendency in human beings. Although cultures (and the people in them) must be responsive to recalcitrant experiences, although human beings are forwardlooking agents, and although certain lines of development are foreclosed by past development, cultures need not develop in a foreordained way. Instead, cultures and the people within them respond to the problems they face (problems that may involve much more than mere survival) based on the situation they find before them, their existing cultural tools, and the available sources of variation or innovation. Although cultures seem to evolve away from the past (albeit at different rates), it is quite unclear what they are evolving toward. It remains entirely possible that human beings will destroy themselves through culture, or reach cultural dead ends and blind alleys of cultural development.

The last chapter offered some analogies between the cultural bricolage of human beings and the evolutionary design of natural selection without suggesting that the two are fundamentally the same kind of process. In this chapter, I want to take up the question of cultural evolution more directly. I argue that there is a significant Darwinian mechanism at work in cultural evolution. However, it does not operate through the natural selection of human beings or groups of human beings. What is replicated and selected in cultural evolution is not human beings but cultural information and cultural know-how in human beings. What is replicated and selected in cultural evolution is cultural software.

Evolution by natural selection requires the "differential survival of replicating entities" in a given environment. More specifically, it requires (1) entities that replicate, (2) a source or mechanism of variation that continuously provides differences among entities, (3) a means by which variations can be passed on to future replicants, (4) an environment in which the entities replicate, and (5) different degrees of survival for different entities within the environment. If all five conditions are met, a process of natural selection results, producing highly complex and differentiated entities over time.2

Nothing in this formulation requires that the replicating entities be organic in nature; the first self-replicating entities on this planet may even have been bits of clay, whose slower replication was swamped by the earliest forms of organic life.3 Hence the principle of natural selection should also apply to units of cultural know-how.

Memes and the Evolution of Cultural Software

Richard Dawkins has coined the word meme (rhymes with cream) to describe these units of cultural transmission. Meme derives from the Greek mimesis, or imitation, and may also be considered to be a pun on the English memory and the French même (same).4

Memes are the building blocks of the cultural software that forms our apparatus of understanding. Memes are spread from person to person by observation and social learning-either face to face or through media of communication like writing, television, or the Internet. Through observation and social learning, people internalize and assimilate skills, beliefs, attitudes, and values, and these become part of their cultural software. In this way, memes are communicated from mind to mind, are adapted into our cultural software, and become a part of us. Culture is a system of inheritance: we inherit our cultural software from the people around us, and we pass it on to those whom we in turn communicate with.5

We use memes to understand, yet memes also "use" us, because they are inside us. Our tools of understanding are constructed from and with the skills and abilities that memes collectively provide. A person is a human being inhabited by memes, a complicated symbiosis of organism and cultural skills. People are complex combinations of their biological inheritance and cultural software, mediated through environmental influences; the information they carry is a combination of their genes and memes.

There are as many different kinds of memes as there are things that can be transmitted culturally. They include skills, norms, ideas, beliefs, attitudes, values, and other forms of information. Examples of memes (or groups of memes) include how to perform a particular dance step; how to build a flying buttress; a tune; a political slogan; how to order a meal in a restaurant; and belief in a divinity. Memes are primarily skills and abilities, but they also include beliefs about the world, paradigms of research, expectations about appropriate conduct (including the conduct of others), lyrics to songs, and ways of pronouncing particular words. Memes encompass all the forms of cultural knowhow that can be passed to others through the various forms of imitation and communication.

Linguistic abilities are primary examples of memes, but so, too, are bodily or kinesthetic skills, for bodily movements are as important to culture as belief systems. Body language and dance; athletic, artistic, and craft skills; gestures, expressions, and other bodily movements—all are to some extent transmissible and hence can constitute memes or complexes of memes. Indeed, imitating and improvising bodily movements may be one of the most basic forms of cultural transmission.

Most writers on the subject have thought of memes primarily as beliefs or ideas that can be stated in propositional form.⁶ This equation is unfortunate. Transmission of culture is primarily transmission of cultural know-how. That is one point of the tool metaphor; a tool allows the self to do something.⁷

Much of the richness of cultural life is lost when we insist on reducing skills to information of a propositional form. Culture does involve information that can be stated in propositional form. Yet this information is valuable because it enables. Hence in describing units of cultural transmission, we must understand knowing that, or even believing that, as a special case of knowing how.⁸

Other writers have argued that representations are the basic units of cultural transmission. Dan Sperber, for example, distinguishes two basic types of representations. Mental representations are beliefs, intentions, and preferences. Public representations are signals, utterances, texts, and pictures; they include what other writers have called symbolic forms. Sperber's emphasis on representations, while helpful, is also incomplete. It does not take into account cognitive mechanisms like associations. As Sperber points out, the most important fact about public representations is that they represent something to someone. The question left unanswered is what allows them to have this representative character. The answer must be in terms of certain cognitive skills that have also been transmitted to others. These skills cannot be reduced to either beliefs, intentions, or preferences. Hence in addition to representations, a theory of cultural transmission needs to grant a central place to cognitive skills.

Although beliefs and mental representations are surely part of cultural software, they are not the whole story. We are more than collections of or receptacles for beliefs and representations; we are embodiments of cognitive skills that produce and interpret beliefs and representations. Focusing on the centrality of cognitive skills helps us remember that culture enables the mind rather than simply fills it up.

The standard view of memes as beliefs is remarkably similar to the standard view of ideology as a collection of beliefs. Both conceptions are unduly limited. Understanding ideological phenomena requires us to look at psychological and cognitive mechanisms that produce beliefs. They include informational filters, heuristics, narratives, scripts, associations of meaning, and metaphoric and metonymic models. These mechanisms are also culturally transmitted and are endemic to cultures and the beliefs of their members. These forms of cultural software are the major concern of Chapters 8 through 11 of this book.

Our ability to assimilate new cultural software often involves the use of existing cultural know-how and hence employs memes or complexes of memes that have previously been transmitted and internalized. In order to learn a theorem in physics, for example, a person must already be able to speak a language, must already have some knowledge of mathematics, and so on.

Memes prepare the way for the absorption of other memes—this progression is another example of the cumulative nature of conceptual bricolage.

The theory of cultural software distinguishes between a person's apparatus of understanding and the public representations or symbolic forms that people understand. Cultural software properly refers to the former and not to the latter. Cultural software consists of tools of understanding that exist within minds. Units of cultural transmission can be stored outside minds—in writings or computer disks, for example—and they can be manifested in utterances and public symbols. But these forms of storage and these manifestations are not cultural software. Nevertheless, the information contained within them can become part of a person's cultural software when it is understood or assimilated through communication or social learning. Although a book does not, strictly speaking, contain cultural software, reading a book may add to or alter our cultural software because we absorb or are influenced by information contained in the book.

This distinction is important because the word meme has generally been used more loosely to describe both units of cultural transmission that exist outside of a person's apparatus of understanding and units of cultural transmission that have become elements of that understanding. Thus we can say that a book or a television program contains memes, and that people absorb memes from watching television or reading books. Memes absorbed in this way can then become part of a person's cultural software. So not all memes are currently part of some person's cultural software, although all cultural software consists of complexes of memes that have been assimilated into or initially created by minds.

Memes, like genes, are units of inheritance, but the inheritance is a cultural inheritance. We inherit our genes from our parents. But we can inherit our memes from anyone we learn from, imitate, or communicate with. We pass our genes on to our children. But we can pass our memes on to anyone who learns from us, imitates us, or communicates with us.

Evolutionary biologists distinguish between the genetic information coded in genes (the genotype) and the physical or behavioral effects of this coding on an organism in its environment (the phenotype). In the same way, we must distinguish between the information coded in memes (the "memotype") and the cognitive and behavioral effects that the meme produces in a person (the memetic phenotype). But because we do not yet know precisely how the brain stores information, beliefs, and skills, we can say very little about the memotype, and we must study memes largely through studying their phenotypic effects.

Memes must correspond in some way to features of the human brain, but we do not yet know exactly how. Each brain is different and may store information in different places. There may be no uniform way that information is stored in different brains, and hence the comparison to chromosomes may be particularly inapt. People whose brains are damaged can sometimes relearn skills using other parts of their brains. Moreover, as Daniel Dennett points out, it would be amazing if "the brain-cell complex that stored the original meme for bifocals in Benjamin Franklin's brain was the same as, or very similar to, the brain-cell complex that is called upon today to store the meme for bifocals whenever any child in Asia, Africa, or Europe first learns about them—by reading about them, seeing them on television, or noticing them on a parent's nose." Nor can we say that the culturally transmissible skill of cello playing corresponds to particular chemical and physical states in the brain coupled with particular configurations of muscles in the hand and arm that are identical for each individual cello player. What makes two examples of a meme in different persons the same is the similarity of the cultural know-how they provide, not the similarity of the ways they are stored in the human body.

Dennett argues that what is preserved in cultural transmission is cultural information in a media-neutral, language-neutral sense. One need not make this assumption, however. First, as media theory reminds us, the medium of transmission may be an important part of the message conveyed. Second, the idea of a media-neutral content of information presumes that social communication essentially involves coding and decoding an identical message. Yet social learning and communication of bodily skills may in fact be much more complicated than this.

For example, the process of advertising does not involve merely a coding of information that is designed to be decoded. Rather than simply convey information, it tries to create similar preferences in different people. Much human communication requires the parties to infer and supplement what is being conveyed rather than simply uncoding it. Finally, the metaphor of coding and decoding an identical media-neutral message is particularly unhelpful in describing how someone teaches another to kick a football, shape pottery, or play a musical instrument. In such cases, we should rather say that the mind and body, through social learning, create their own individual skills similar to but not necessarily identical to those perceived in others. This is a form of replication, to be sure, but decoding is not the appropriate metaphor.

Genes usually replicate in complexes or groups, called genomes. It is likely that culture is also transmitted in complexes of memes, or memomes. Before the discovery of the biochemical vehicles of genetic inheritance, it was difficult to determine where genes began or ended. Scientists had to make inferences about the boundaries between different genes from their phenotypic effects on an organism's physical features and behavior. Often (as in the case of blue eyes) a phenotypic effect is the result not of a single gene but of a combination of

genes, but scientists could not determine this until they understood the biochemical basis of inheritance. Because we do not know precisely what biological vehicles carry memes, it is hard to separate the meme from the memome in the way that we can now separate some genes from their genomes.

These limitations in our knowledge raise a problem of demarcation and division. Is Beethoven's Fifth Symphony a single meme, or does the four-note motto that begins the piece qualify by itself? The solution to this difficulty is entirely pragmatic, as it is in the case of genes. Multiple traits are often passed together from parent to offspring—for example, a certain shape of nose and a certain eye color—but we can say that the gene is the smallest unit of genetic information that can be and is repeatedly transmitted more or less intact. In a similar spirit, we can say that memes are the smallest units of cultural skills or information "that can replicate themselves with reliability and fecundity." ¹⁴

This solution does not eliminate all difficulties. Daniel Dennett argues that that the notes D-F#-A do not constitute a meme, while the theme from the slow movement of Beethoven's Seventh Symphony is a meme. Just as a single codon of DNA like C-G-A (coding the amino acid arginine) is "too small" to be a gene, Dennett believes that the effects of the notes D-F#-A are insufficiently individual to count as a meme. A "three nucleotide phrase does not count as a gene for the same reason that you can't copyright a three-note phrase: it is not enough to make a melody."15

Dennett's argument misses an important difference between genetic transmission and cultural transmission. The biochemical vehicles of genetic transmission place lower limits on the size of the units of transmission. A three-nucleotide phrase cannot be a gene because of the biological structure of genes. But cultural transmission works very differently. A skill or a piece of information can be a building block of other, larger elements and yet also operate as a meme in its own right in other contexts, as long as it has some independent memorizable meaning to an audience. Moreover, the length of the sequence is not the only factor. The cultural expectations of audiences (which include their own preexisting cultural software) help determine what is reliably memorable and what is not. Thus, the musical phrase F#-E-D played at moderate tempo represents the song "Three Blind Mice" to people living in certain cultures. Because these three notes played slowly call that song to mind, they can serve as a symbol of the entire piece. And a symbol—something that stands for something else to someone in some context—is a particularly salient example of a meme. Note, however, that these three notes played slowly also begin the second subject in the first movement of Tchaikovsky's Pathétique Symphony. 16 The first F#-E-D is a meme, but the second (at least currently) is not.17

We now see why Dennett's analogy to copyright law is mistaken. The

reason why the law does not permit short phrases to be copyrighted has nothing to do with the phrase's inability to have significance or to be memorized reliably and repeatedly. It stems from the fear that the owners of the copyright could demand royalties for each and every use, thus stifling creativity. Protection is denied not because short phrases cannot be memes but because they can be memes. Because units of cultural transmission can act both as independent units of meaning and as building blocks for other units, the law does not allow the very smallest units to become intellectual property; excessive property protections may block larger, socially beneficial constructions.¹⁸

This example raises one of many important differences between biological and cultural evolution. Memetic evolution may be a process of natural selection, but it does not necessarily occur in exactly the same way as biological evolution, or use precisely analogous structures and techniques. Many features of biological evolution may result from the particular requirements of biological replication and designoid structures arising earlier in the development of life on this planet. For example, biological evolution on Earth makes use of DNA and RNA, alleles and codons, because of the particular way that life originally formed and was able to reproduce itself. There may have been many different possible biochemical structures of biological reproduction and evolution, but organisms on this planet hit upon a particular one and successively built upon it. The structures that we discover in genetic evolution may not be in any sense necessary to evolution but may simply be the ones that evolved historically to transmit genetic information—given the constraints of the particular biological organisms that were first able to carry and reproduce this information through their own growth and reproduction. It does not follow that all forms of evolution through natural selection require analogous structures of transmission and evolution.

We should use the concepts of biological evolution to the extent that they can serve as a useful heuristic to understanding cultural evolution. We can start with a model of evolution that we already know something about and use it as our point of departure for studying other forms. But if we rely too heavily on biological analogies, we will inevitably be misled, because biological evolution is only one possible form of evolutionary development. We must always be on the lookout for disanalogies. Indeed, discovering these disanalogies often is as helpful in understanding cultural evolution as discovering analogies.

Memes as Populations

Like genes, memes are self-replicating entities, but the environment in which they replicate consists of human minds and the places for external information storage that humans have devised. At any point in time there is a "meme pool" of memes competing for survival in the environment of human minds, just as there is a gene pool that competes in its environment. No two human beings have the same memes; there are no identical cultural twins. Every human being is a unique individual, and not simply the replication of a cultural template.

Both the study of cultural evolution and the study of genetic evolution by natural selection employ what Ernst Mayr has called "population thinking." Species are populations of slightly different organisms that carry slightly different combinations of genes. The frequency of these genes in the overall population changes depending on how organisms interact with the environment. Individuals with genes favored by their environment have better chances to survive and produce more offspring. Over time, the species evolves because of the changing composition of its gene pool. Viewing species as populations means that species are not essential, unchanging types, and members of a species are not imperfect examples formed from a standard template. "There is no 'typical' individual," Mayr points out, "and mean values are abstractions." Variation is not only characteristic of individuals, it is essential to the forces of change. As Mayr puts it, "he who does not understand the uniqueness of individuals is unable to understand the working of natural selection."19

A similar analysis applies in the world of culture. Each person is constituted by a population of memes—her cultural software—and the entire population of human beings represents an even larger population of memes. We can think of cultures, subcultures, and interpretive communities as populations of partly similar, partly different memes reflecting partly similar, partly different cultural software in individuals. Cultures, subcultures, and interpretive communities are neither natural nor supraindividual entities; they are effects of or useful abstractions from the slightly different cultural software of their members. Cultures have conventions and institutions that help to reproduce the cultural software of their members. Yet these coordinated behaviors are also the effects of the similar but slightly different cultural skills of their members. The cultural and the individual thus tend to fade into each other: what is cultural consists of widely spread and long-lasting memetic features of individual members of the culture, just as the species consists of the widely spread and long-lasting genetic traits of individual members of the species.²⁰

We can reinterpret the concept of a cultural tradition in these terms. There are two ways of understanding a tradition: one synchronic and the other diachronic. These two visions of tradition correspond roughly to two different ways of thinking about species—synchronically, as a population of relatively similar individuals with relatively similar genes existing at a given point in time, and diachronically, as a line of genetic descent.

Viewed synchronically, a tradition is a set of ways of thinking through which people understand and live at a particular moment in their lives. Traditions involve populations of relatively similar memes that constitute the people who live within those traditions. A tradition produces a particular perspective or horizon—the way of seeing and understanding the world that makes use of the cultural software that one has. The shared perspective among the members of the tradition is due to the similarity of their tools of understanding.

Diachronically, a tradition is an ongoing process in which members inhabit one particular, though not necessarily privileged, moment. The tradition is reproduced in successive generations, but it changes over time and may have transformed itself considerably in the process. What allows people to think of themselves as "traditional" Jews, for example, is not the belief that they are doing exactly what people did 3,000 years ago in ancient Israel. Rather, it is the belief that there is a genealogical continuity between what others who called themselves Jews did in the past and what one is doing now. This approach views tradition as a line of memetic descent, in which the memes possessed by the members of the present interpretive community can be seen as linked through a chain of communication and education with the memes of earlier members. Because memetic evolution occurs much more rapidly than genetic evolution, traditions may evolve and change quickly, and over time many of their core beliefs, practices, and rituals may be displaced, despite our ability to trace a transhistorical continuity of transmission.

The environment for memes consists of human minds and methods of memory storage. There is a limited number of minds in a geographical area, in a particular culture, or in the world. Each mind has limited time for social learning and limited information-storage capacity. Memorization or achievement of a skill not only requires exposure to cultural transmission; it also requires conversion from short-term memory to long-term memory. Repeated exposure and practice may be necessary if the skill is to endure and become second nature.²²

Moreover, even though memes can eventually be stored outside of minds, they still need the intervention of minds at crucial points for their replication and continued survival.²³ Hence human minds create a bottleneck for the replication and storage of cultural skills and information. Memes must compete for available space. Variation among memes causes different rates of survival and propagation. Memetic competition for available space in the minds of human beings creates gradual changes in the population of memes in a particular geographical area, or in a particular culture or subculture. If the survival rates differ enough, particular skills may become extinct or die out. Languages, for example, die out when insufficient numbers of speakers are available to reproduce them. The same is true with bits of information or bodily skills. If everyone forgets how to do the rhumba (and no external records of how to perform

it remain), the rhumba goes out of existence, although something like it may be invented in the future.

Indeed, not only is there competition among memes and meme complexes for instantiation in human minds, there is competition within each human mind for those memes that are most easily remembered, repeatedly transmitted, and frequently employed. People use some skills more frequently than others and forget some things more easily than others. They bring some skills to bear more often and other skills less often. They develop some skills more and other skills less. They think and talk about some things more frequently and others less so. We might say, loosely speaking, that a human mind is a population of competing memes that exists in a larger population of competing memes called a culture, a subculture, or an interpretive community.

One of the most important parts of the environment that memes face are other memes in the meme pool and the behaviors and beliefs they produce. Thus memes not only compete, but they also must adapt to the existence of other memes and may even benefit from or depend on other memes for their continued survival. Many memes in human culture survive and propagate only because human beings already have internalized and mastered certain linguistic skills and vast quantities of information and cultural know-how. Previous education is often necessary to comprehend, recall and utilize newer skills and information. Some memes, like some genes, can even be coadaptive, so that they mutually assist in each other's survival.²⁴

Memetic Variation

An important difference between cultural and genetic evolution concerns the frequency of variation. Genes usually make very good copies of themselves; mutations are a relatively rare occurrence. The same is not true in the world of culture. Cultural transmission requires communication, imitation, or some other form of social learning. The copies produced by this process are rarely identical to the original. Misunderstandings occur, or, more frequently, partial understandings occur that are good enough for one purpose but not for another, unforeseen purpose. Skills require practice to be perfected; the need to practice them means that the earlier attempts will be inartful and that later attempts will draw heavily on the recipients' own personality and abilities. Just as no two people dance, cook, or play the cello in precisely the same way, no two people understand social conventions or situations the same way. As they pass these skills and understandings onto others, further change occurs.²⁵

Memetic mutuation occurs not only because of misunderstanding or because old signs are inserted into new contexts. It also occurs because of innovation. Human beings are not passive receptors of memes; they are active processors and recombiners of the cultural messages and skills they receive from others.

First, human minds combine and adjust the memes they receive with those they already possess.²⁶ For example, people have filtering mechanisms for receiving new information; they may discount information if it conflicts too greatly with what they already believe.²⁷ Similarly, the theory of cognitive dissonance suggests that people may reconfigure new ideas and understandings to achieve intellectual coherence with their existing beliefs or to preserve their sense of themselves.²⁸ People also create new memes when they learn through trial and error.

Second, individuals are creative. They modify skills, combine information, draw inferences, and stretch conventions. To be sure, people always do these things by making use of the cultural software they already possess. But this fact does not make their activity any less creative; indeed, their cultural software enables their creativity by providing thought with a necessary framework for problem solving and innovation. In short, human beings are not simply Xerox machines for their memes; they are also incubators for new memes, as well as master chefs who combine old memes to create new memetic recipes. We send our newly created memes out into the world, where they are received, assimilated, adjusted, recombined, and modified by countless other minds, each creative like our own. The power of human reason, made possible in part by the memes we possess, is also the power to mutate those memes and create something new from something old. We are not simply the inheritors of a zealously guarded patrimony but entrepreneurial producers of new cultural software, which will help constitute future generations of human beings. So the story of memetic evolution is neither the story of our slavery to memes nor the story of how human reason enables us to break free of this slavery. Rather, it is a story of the collective creation of human reason, a story of powers of heightened creativity made possible by previous memetic infestations, a story of freedom mixed with, and paradoxically made possible by, constraint.

Because human beings are creative and combinatory, the path of cultural evolution must necessarily be different from that of genetic evolution. In the Origin of Species, Darwin used the metaphor of the branches of a tree to describe the basic trajectory of evolution.²⁹ Life on Earth, he argued, has a single origin. Different species diverge from this root at different times, further subdividing into new species. This topology means that as species evolve, they separate into ever new forms, and the proliferating branches never recombine. But the history of cultural development is quite different. Cultures do tend to diverge because of geographic isolation or disciplinary specialization, but later people often borrow from other cultures to supplement their own. For example, American cuisine—which includes such delicacies as deep-dish pizza, hamburgers, and chop suey-borrowed from other cultures to create dishes that did not originally appear in those cultures. (Though to be sure, the spread of American culture in the twentieth century means that now one may indeed be able to get hamburgers in Hamburg, deep-dish pizza in Rome, and chop suey in Hong Kong—yet another example of cultural recombination.) A similar point applies to technological development. Inventors often look for solutions by lifting ideas from widely divergent cultural sources: the use of computer punch cards to store information, for example, was inspired by the Jacquard loom, which was in turn based on the earlier technology for constructing automated pipe organs.³⁰ This sort of cultural borrowing is yet another example of the bricolage described earlier. In sum, one of the most important distinctions between genetic and cultural evolution is that while biological lineages increasingly diverge, cultural lineages often recombine.31

A second important distinction concerns the mechanisms of replication and survival. As noted earlier, in the cultural world, transformation is the rule and exact copying is the exception. Moreover, much cultural transmission is not a process of coding and decoding an identical message; it may involve creating similar cognitive skills through imitation and inference from salient examples. An evolutionary theory of culture based on the differential survival of replicating entities must take these facts into account. If memes are constantly being transformed as they spread, the mechanism of differential survival must operate differently in the cultural and biological worlds.

The survival of cultural software does not depend solely on different rates of attractiveness or acceptance by human minds. Human beings inevitably transform what they receive from others; even if I like what I see or hear, it will be changed when I pass it on to others. This presents a real problem for memetic survival: if transformations were purely random, they would eventually destroy the identity of what spreads. If people randomly transformed different elements of an original story each time it was told, after a time there would be not a single version that was widespread but a random distribution of many different stories. For memes to be successful replicators, it is not enough that they have descendants; they must also have sufficiently similar descendants.

A particular kind of cultural software will not become widespread in a population unless its transformations are systematically biased in particular directions, or tend to converge on a central set of features. Put another way, if a particular kind of cultural software does become widespread, it is probably because some types of creative change or transformation of that software are more likely than others. What becomes widespread is not only what resists transformation but what gets transformed, in roughly the same way by many different people.32

Human psychology and cultural factors play a central role in these subsequent transformations. For example, people are more likely to retell those elements of a story that are most salient and easily memorized; they are more likely to forget or transform others. In this way psychological properties of memory and relevance determine how some parts of stories are retained, how other parts get transformed, and how those transformations converge. Transformations may also converge because of the force of existing cultural expectations. A story with an inconclusive ending will probably gain a happy or a sad ending if it is repeated often enough. Here multiple transformations may lead to the spread of two different stories with different endings. Conversely, the process of transformation may tend to combine different varieties of cultural software. Two or more different stories may eventually converge into a single story because successive recountings of each get transformed toward a common version.

In describing the spread and success of cultural software, therefore, we must consider effects on the "demand" side (what kinds of memes are most attractive, salient or useful to other minds) and on the "supply" side (what kinds of transformations memes will undergo as they are communicated to other minds). In the biological world, the problem of guaranteeing similarity among descendants is not very great because the biochemical mechanisms of copying tend to be fairly accurate. Hence the most important determinant of reproductive success is selection by the outside environment. But in the cultural world, a meme must contend with both the outer environment of other minds that might be receptive to it, and the inner environment of the mind that propagates and transforms the meme. It must survive in both environments, and it must survive in ways that retain its commonality with other memes.

Problems of Transmission

In order for memes to replicate, they must be embodied in some vehicle. People are the most important vehicles for memes, but books, records, and computer disks also serve as vehicles for cultural replication. Technology itself can serve as a meme vehicle. The very existence of a wheel suggests to us the fact that such a tool can be created, how to create it, and how to use it.³³ The amount of information technologies convey about themselves is necessarily limited, however, especially as the technologies get more complex.

If all of a meme's physical embodiments (including all human memory storage) are destroyed, the meme becomes extinct, although something like it can be invented anew. The durability of a particular vehicle does not necessarily guarantee reproductive success over time. Many insect species have existed for millions of years, even though the lives of individual insects are comparatively

short. Rather, it is more important to produce many copies than to ensure that all the copies survive for long periods of time.³⁴

Cultural evolution is not possible until there are sufficiently powerful information-processing devices capable of storing information and reliably transmitting it to or replicating it in other information-processing devices. Animals have rudimentary abilities to produce culture and pass it along to their offspring. Animals can learn skills and imitate movements, and some even have rudimentary semiotic and linguistic skills.³⁵ Birds can imitate songs and transmit them from generation to generation, and these songs can even mutate over time.³⁶ But if a pigeon sees a copy of Shakespeare's *Hamlet*, it cannot assimilate the memes contained in that play. Even if an animal could memorize a particular skill or particular information, the skill or information dies with the animal if it cannot transmit its mastery reliably to others. Fecundity, transmissibility, and longevity—three essential requirements for a process of natural selection were not sufficiently present before the evolution of humankind.

Once comparatively large-brained human beings arrived on the scene, and invented language, however, memetic evolution really took off.³⁷ Fecundity was greatly increased because many different kinds of memes could be transmitted to many people at once through vocal communication, observation, and imitation. The memes involved in linguistic ability greatly enabled the transmission, processing, and storage of other memes, which in turn enabled the transmission, processing, and storage of still others. Longevity was enhanced because even though a particular person died, her information could be passed on to others. This environment was still somewhat inhospitable for memes because it relied so heavily on human memory for storage and on human speech and movement for transmission. The next great advance in memetic fecundity, transmissibility, and longevity was the invention of external forms of information storage: first through writing, then by means of printing presses, and in our own day through the use of digital computers. With the invention of writing it became possible for the ideas of an ancient scholar like Plato to survive into this century without having to be fully memorized by an unbroken chain of individual memories. Indeed, to the extent that external forms of information storage are more durable than human memory, there may be some comparative advantage for memes to convert human memories into these more durable forms. It is this possibility that underlies Dennett's wry suggestion that a "scholar is just a library's way of making another library."38

The development of extrinsic sources of information storage is important for another reason. As we have seen, the human mind is a natural bottleneck for memetic evolution, because memes usually must reside in a human mind before they can be transmitted to others. The scarcity of human minds is an important element of the natural selection of memes. Increase in the brute number of human beings eases this bottleneck to some degree: for example, it makes possible (but by no means guarantees) a flowering of knowledge that might not have been possible in earlier times. Nevertheless, the bottleneck remains.

On the other hand, if computers become sufficiently developed, human mediation and incubation of memes may become increasingly unnecessary. To begin with, after a certain level of technological innovation is reached, it becomes possible for information to propagate without its contents being directly stored or understood by any human mind. Whenever computers communicate with each other or copy files, for example, information is propagated whether or not it is ever accessed by a human mind. A human mind is still necessary to design, program, and repair the computers, but it is not necessary for a human mind to think all the information the computers contain. Eventually, it is possible that more and more features of maintenance, programming, and design could be left to computers themselves. A rudimentary example is the current use of computers to design computer chips.

The creation of new propagation and incubation devices might ease the bottleneck of memetic growth and thus drastically change the course of memetic evolution. Indeed, it is quite possible that some memes may presently find computers a more hospitable environment for development than the human minds that their ancestors originally inhabited (and spurred on to construct computers). Of course, the features that benefit a meme's survival and propagation in a computer's memory banks may be quite different from what guarantees its survival in a human brain or on a piece of paper. Hence memes that successfully inhabit computers may evolve differently and possess somewhat different features from those of their human-dwelling cousins.

In order to reproduce successfully, memes must be able to transmit themselves from one mind to another. Originally, this must have presented an enormous hurdle. One cannot simply copy cultural software onto a brain as one would load software from one IBM-compatible computer to another. Copying software is easy on these computers because each has identical physical structures for reading and coding data and an identical hard-wired machine language. By contrast, copying and running software on computers with different and proprietary hardware is actually a fairly difficult task. At the beginning of the personal computer revolution in the 1970s and early 1980s, for example, there were literally dozens of incompatible computer designs, none of which could load or run one another's software. Many computer manufacturers eventually went bankrupt because their machines were not 100 percent IBM compatible. Our image of computer software as something that can be easily popped out of one computer and into another is really the result of competitive pressures that weeded out most designs for personal computers until the IBM and Apple designs achieved market dominance in the mid-1980s.

Human beings are decidedly not like mass-produced IBM computers. The physical structure of each person's brain is different, a product of both genetic inheritance and subsequent development. The mental capabilities of human beings are more like a proliferation of different proprietary architectures, each with its own unique features. If we want different kinds of computers to talk to each other, we must create a program on each that can accommodate its architectural idiosyncracies. Each such program creates on the computer a "virtual machine" of software and hardware that can read and understand common instructions, and thus can speak a common language.³⁹

Because each brain's structured capacities are different, memetic exchange must occur through a mode of transmission that is, as Daniel Dennett puts it, "social, highly context-sensitive, and to some degree self-organizing and selfcorrecting." Put another way, if human beings can transmit and share cultural software, it must be due to the differential survival of memes that have a high degree of adaptability and tolerance for different mental environments. 40 At the same time, there must be some degree of commonality in the basic cognitive and linguistic apparatus of human beings to allow such hardy memeskills to have developed in the first place. The scope and extent of this universality is the well-traveled terrain of the debate between Chomskyites and their opponents.

These transmission skills come in several varieties, including learning by imitation, through positive or negative reinforcement, and through natural language. 41 Once these skills exist even at the most rudimentary levels in brains that are big enough for large numbers of memes to inhabit, the process of memetic evolution takes off, building its own "information superhighway" from previous meme-skills and facilitating the replication of more and more memes. Thus, just as human beings change their environment to make it more hospitable, memes without intention or plan develop and combine to create a more hospitable environment for themselves both in human brains and in extrahuman information-storage devices. The development of the first hardy memes that could create the virtual machines that facilitated transmission was itself the result of a process of natural selection. Memes that were able to do so successfully spread, while those that could not failed to take hold in the meme pool. Subsequently, other memes could and did take advantage of this newly fertile ground.

Memes as Filters

As I noted earlier, most theorists who discuss units of cultural transmission have focused on meme-beliefs, rather than meme-skills. This bias is similar to the general tendency to assimilate all features of ideologies to beliefs. Yet many of the most important forms of cultural software—and particularly the most important for the study of ideological phenomena—are skills or cognitive structuring mechanisms that cannot be reduced to propositional beliefs. A good example of such a meme-skill is a filtering mechanism.

It is clear that beliefs can act as filters; an example is the notion that one shouldn't believe anything printed in a particular publication or spoken by a particular politician. ⁴² But filters do not have to exist in the form of propositional beliefs. Many cognitive mechanisms, including prejudices, narrative structures, metaphoric models, and metonymic associations, act like filters. They let in ideas that conform to particular patterns of thought while rejecting those that do not. Psychologists have also discovered a series of heuristic mechanisms that people use to search for information and other mechanisms that people use to assess and discount information contrary to what they already believe. ⁴³ These mechanisms also filter experience. Alternatively, cognitive mechanisms can actively adjust and shape new social experience so that it appears to conform to existing structures of thought and belief. Mechanisms of cognitive-dissonance reduction seem to work in this way. ⁴⁴

We can put this point more generally. Many ideological effects are produced by memes that act as cognitive filters. There are many different ways that our cultural software can do this, and the study of how it does so is a large part of the study of ideological effects. Memes that act as cognitive filters become part of the environment for new memes that seek entry into human minds. These filtering memes help the mind to accept some meme candidates and reject others, or help adjust and reconfigure incoming memes to existing patterns of thought. Hence these meme filters are part of the mechanisms of natural selection that occur within each individual human mind. All other things being equal, memes that can most easily break through or accommodate themselves to the filtering mechanisms of an individual human mind are more likely to find room in the limited memory space available in that mind.

Moreover, because filtering memes help determine which memes are accepted in human minds and which are not, they are important mechanisms of natural selection of beliefs and skills within cultures, and indeed, across the entire population of human minds. All other things being equal, those memes that can most easily break through or accommodate themselves to the filtering mechanisms of people's minds will, over time, be more represented in the meme pool of a given population, culture, or subculture.

In sum, meme-filters help explain how human beliefs—and hence the ideological effects of those beliefs—develop and spread differentially. The idea of a meme-as-cognitive-filter links the study of memetics or cultural evolution to the theory of ideology.

But if meme-filters are an important source of ideological effects, and if

they are part and parcel of the natural selection of memes in the ecology of human minds, how are they themselves selected for in the first place? Why would human minds develop meme-filters as part of their cultural software?

The most important fact about information is that there is too much of it. Finite human minds need ways of taming the Heraclitean flux of experience. Thus filtering, organizing, and structuring information is a positive good, and memes that act as filters naturally arise to fill this need. Like all evolutionary innovations, such filtering mechanisms do not have to be perfectly designed. They need only be good enough for the purpose at hand and may have all sorts of unforeseen and unforeseeable side effects. This means that some mechanisms of filtering may be harmless or even quite helpful in some contexts but harmful, prejudicial, and unjust in others.

Our memes filter, organize, and structure social experience. They provide key components of the environment in which new memes will grow, develop, propagate, and perish. These filters and structures arise along with the proliferation of information. Hence increasing the amount of available information does not necessarily increase knowledge or understanding. It does not result in a person's being well rounded or well read, having an open mind or being receptive to new ideas. Indeed, under certain circumstances it can have precisely the opposite effects.

Encountering an explosion of information can foster closed-mindedness, because too many competing sources of information produce the potential for confusion. The flood of conflicting information creates a suitable environment for breeding ever new forms of memetic filters that harness the flow and shut out many different kinds of information. Some of these filters may include mechanisms that hide their own biases and limitations, because this tends to increase their success at propagation.

Thus we should not necessarily assume that the proliferation of new information sources and the coming together of many different cultures will produce the end of ideological conflict. Rather, the widespread availability of information and the collision of many different cultures and language games may in fact produce more narrowness of thinking, more inflexibility, and more intolerance, whether between ethnicities or between academic disciplines.

The development of memetic filters creates new bottlenecks for the propagation of memes. Memes that can break through or get around these filters have greater chances to spread in a population of minds. Thus complexes of memes develop means of evading filters: examples are the development of flashy graphics, large type, or loud music in advertising to attract an audience's attention. In response, new forms of filtering arise to keep pace. The result is a sort of "arms race" between memes seeking places in human minds and the filters designed to winnow them out.⁴⁵

Eventually, filtering and devices to get around these filters start to exist in symbiosis, so that it becomes difficult to distinguish between what is filtering and what is promotion designed to evade filtering. Advertising the status of an author, for example (through institutional affiliation or kudos on the dust jacket), can be used to signal that the work is worth reading, but this signal is also adjusted to known filtering mechanisms for deciding which works to read. Signals, in other words, are the flip side of filters. They are devices used to advertise quality or desirability to a potential audience so as get past informational filters. And as Dennett observes, "'Blind refereeing,' the proliferation of specialized journals, book reviews, reviews of book reviews, and compilations of 'classic works'" can be seen both as filtering devices and as means to get through these devices and into human minds.⁴⁶

Memes as Viruses

Because cultural software is transmitted from person to person, there is a natural analogy between cultural software and viruses. The human mind is susceptible to memes just as the human body is susceptible to infection from particular viruses.⁴⁷ The study of cultural evolution is a study of comparative epidemiology. Some memes are more contagious, or "catching," than others in a population and thus spread more widely and successfully.⁴⁸

The metaphor of susceptibility to viruses helps us understand the deep connections between the power of human intelligence and its vulnerabilities. Human beings are more susceptible to many more kinds of memes than, say, pigeons, precisely because they have a greater intelligence. All forms of human understanding involve susceptibility to memetic invasion. Human beings are vulnerable to memetic infection precisely because they are so well developed as meme reception machines. Moreover, much of this infection does not involve someone intentionally sending a message to a recipient. We receive memes when we observe another person's behavior or dress or when children pick up ideas or behaviors from their parents, teachers, or schoolmates. Just as children easily contract all sorts of diseases, they are particularly susceptible to memetic "infections" in all sorts of unintended ways. That is one reason why parents are so particular about what their children are exposed to.

People's susceptibility to memes varies with the skills that they already possess. Our existing cultural software shapes what is salient, interesting, and hence what is easily communicated and easily absorbed. Although children learn all sorts of words that make their parents cringe, they are relatively immune from discussions of fluid dynamics. I am susceptible to memes in English but relatively immune to those in Urdu, because I do not speak that language. Lawyers who practice bankruptcy law are more susceptible than

laypersons to remembering and being affected by discussions of recent bankruptcy decisions.

Taking the metaphor of disease one step further, we can describe a continuum between two types of cultural infections, the endemic and the epidemic. Some cultural software is transmitted persistently over generations of individuals and through many different human cultures; it is endemic to a particular culture or to human thought generally. Other forms of cultural software spread rapidly from person to person, like advertising slogans and fashions. These memes are comparable to epidemics.50

The distinction between epidemic and endemic cultural software is quite important for the study of ideology. Many of the cognitive mechanisms that produce ideological effects in human thought are endemic rather than epidemic. Narratives, networks of association, metaphors, and metonymic models are transmitted widely and persistently. Moreover, once in place, these endemic forms of cultural software provide the environment in which epidemic cognitive structures and beliefs can thrive. The study of ideology is the study both of endemic cognitive structures and of epidemic changes in beliefs and symbols.

Racist thought can be both endemic and epidemic. Racist thinking occasionally sweeps from person to person like a dangerous virus. Yet equally important to understanding the phenomenon of racism are more basic cognitive structures-for example, historical associations of white and black with contrasting positive and negative stereotypes. These networks of association are endemic—they are deeply embedded and widely reproduced in many cultures. They prepare the ground for the development and spread of racist beliefs.

Memes as Symbionts

The account of cultural evolution that I have been developing suggests that not only do people have ideas, but ideas have people. Memes "use" people for the purpose of their own propagation. We should not understand such anthropomorphic language literally: memes no more than genes have wants, desires, purposes, or interests. Rather, this is merely a shorthand way of describing how natural selection works on units of cultural transmission.

This approach removes the need to explain human cultural development and proliferation solely in terms of its survival advantages for human beings. To the contrary, we may assume that much cultural development is largely irrelevant to human survival in the short term, although it may have many profound and unexpected long-term effects. Memes do not necessarily reproduce and propagate because this process confers an evolutionary advantage on human beings (although this may in fact occur). Rather, they survive, reproduce, and propagate because it advantages them.

Thus, we can think about cultural software as a kind of symbiont. A symbiont is an organism that lives inside or attached to another organism. The latter organism is called the host, the symbiont is called the guest. We can think of cultural software as a kind of informational symbiont. Under this analogy, our bodies (and our brains) are the biological hosts for cultural software.

Roughly speaking, we can divide symbiosis into three kinds. In the first case, mutualism, the host and guest enhance each other's reproductive fitness. A second case, commensalism, occurs when the symbiosis benefits the guest's reproductive fitness with little or no cost to the host. The third and most familiar case is parasitism. A parasite is a guest that benefits at the expense of the reproductive fitness of the host.⁵¹ The natural world does not divide up as neatly as these categories would suggest. It is quite possible for a guest to help the host in some ways, harm it in others, and be neutral in still others.

Memes are like symbionts that alter the behavior of their hosts, much as the rabies virus alters the behavior of a dog by making it more aggressive, increasing its salivation, and preventing it from swallowing.⁵² Just as the genes in the rabies virus make use of the host to spread their genetic information, memes use their hosts to spread their own memetic information. The rabies virus is a parasite because it increases its own reproductive success at the expense of the dog's. However, the survival and spread of memes can either be advantageous, indifferent, or in opposition to the reproductive fitness of the host—memes can be mutualist, commensalist, or parasitic.

There are two important limitations to the comparison between memes and biological symbionts. First, the union of biological capabilities and cultural software creates a new kind of entity, a person. This is not true in the case of a dog infected with a rabies virus. It is still a dog. Second, this new entity, the person, has new interests independent of the reproductive success of the biological host. People have interests in both senses of that word: there are things that they are interested in (that is, they have preferences, desires, and values) and things that are in their interest. These categories can be further divided into long- and short-term interests, and the various kinds of interests can conflict with each other.

Most living things have relatively uncomplicated interests in eating, surviving, and reproducing. Because people combine genes and memes, their existence is more complicated. Their interests constantly develop, change, and conflict during the course of their lives, and they often have no idea precisely what they want or exactly what actions they should take. Indeed, we might define a person as an entity that is continually at a loss for what to do.

We must therefore distinguish between what advantages a person's interests (in the various senses of that word) and what advantages the reproductive success of his or her genes. Memes can help one while hurting the other, and vice versa. They can be parasites with respect to reproductive success but commensals or even mutualists with respect to a person's other interests. Consider a Catholic priest who takes a vow of celibacy. Reading certain books or mastering certain cultural skills that help him keep his vow would not enhance his reproductive success, but it might further his other interests.

What complicates matters is that some of people's existing cultural software helps to shape and constitute their interests, and this helps determine what is harmful or helpful. For example, memes that lead a person to watch a lot of television may be mutualist for a person who is a television critic but commensal or even parasitic for a person who is a law student.⁵³ In addition, it is often difficult to separate the interests of memes from the interests of the persons whom they constitute. Consider the question of whether being a television critic is really in my interest. Important features of our personality and important choices in our lives may be the result of the cultural software we possess; they may be inextricably linked to our personal identities and our sense of ourselves.

It is likely that the earliest human memes were predominantly mutualists that enhanced our reproductive success.⁵⁴ They helped human beings (and human genes) do things that helped them survive and reproduce. Primitive systems of communication and cooperation may have been the earliest examples of widely transmitted cultural software among humans. They provided the basic skills necessary for social learning and the spread of culture; and they themselves spread because these skills improved human beings' chances at survival.55 The earliest memes probably built on innate skills. Cooperation skills built on whatever instincts for social coordination human beings already had; natural languages built on innate linguistic skills.

But once the first hardy memes took root and spread widely, they prepared the way for other memes that could not previously have infected their human hosts. They created an environment in which new memes could flourish that did not necessarily assist human reproductive success, or that even undermined it. In this way, memes, which originally gained a foothold in human minds as a way of increasing genetic fitness, took on a life of their own. They created new structures for processing information, and thus new susceptibilities for infection by ever more exotic forms of memes, including many commensals and parasites. The new cultural environment in turn created new human interests and hence the possibility of ever new forms of mutualists, commensals, and parasites. As a result, the cultural world we inhabit today contains all three kinds of cultural symbionts.

To be sure, memes are at a severe competitive disadvantage if they routinely threaten the survival of their human carriers; an example would be a belief in the necessity of suicide.⁵⁶ At first one might think this to be true of beliefs that

encourage violent confrontations, war, and murder. But as long as a meme can propagate and reproduce fast enough in enough human beings, the destruction of large numbers of belief carriers is not necessarily fatal to the meme's continued survival. Indeed, to the extent that violence reinforces the violent or aggressive beliefs of the surviving human carriers—for example, by confirming that hatred of the other is justified and that only strength can ensure safety this may even have a salutary effect on the propagation and survival of violent or aggressive belief.⁵⁷ (Compare the rabies virus, which eventually kills the dog but in the process spreads itself by promoting aggressive behavior.) A similar point applies to the many different cultural skills involved in warfare and destruction. As long as the skills involved in killing people do not completely exterminate the earth's population, these skills will find a welcome home in human minds and propagate accordingly.

The complicated relationship between the interests of memes and their human carriers has a partial analogy in biological evolution. Darwinian processes of natural selection can simultaneously occur at several different levels, with the result that they "leak" into each other or have feedback effects on each other.⁵⁸ That is because an entity that is the object of natural selection within a particular environment can also itself be the environment in which another Darwinian process occurs. Human beings are objects of natural selection in their environment, but human bodies and human cells are also environments where lower-level processes of natural selection can occur. Human cells use DNA for replication, but only a small percentage of human DNA actually is involved in providing the necessary codes for constructing proteins. Much of the rest contains sequences randomly dispersed and repeated over and over again, with no apparent function.⁵⁹ Although some of this DNA may indeed have beneficial effects, the best explanation for it lies elsewhere. Simply put, this DNA has found a way to make copies of itself within the "environment" of human cells and does so because of a familiar Darwinian logic: DNA that does not reproduce itself in this way, or does so less efficiently, will, over time, be increasingly less represented in human cells.

Nevertheless, the human body can be adversely affected by the proliferation of these unnecessary copies in human chromosomes. If this repetitious DNA were to completely take over human cells, it would kill them or so exhaust their energies that it would cause a significant disadvantage for the survival of its environment, the human body, and thus pose a significant threat to its own survival.⁶⁰ On the other hand, if the replication of redundant DNA does not significantly damage or otherwise reduce the survival and reproduction of human bodies, then it will not be weeded out by the higher-level Darwinian process. Hence, redundant DNA acts as a sort of "intelligent parasite," multiplying as much as it can, but not so much that it kills the goose that lays the

golden eggs.⁶¹ Thus, there is a sort of feedback effect between the two different levels, an interaction between two Darwinian processes. The lower-level units of selection have an interest, albeit an attenuated and imprecise one, in the survival and propagation of the higher-level units of selection that constitute their environment.

No doubt a similar feedback between the survival interests of human beings and memes is also at work in cultural evolution. As long as memetic evolution has no immediate disadvantage for human survival, it is free to develop in many different directions, with unpredictable long-term effects for the human carriers of culture. Moreover, for those who are suitably paranoid, there may even come a time when computers and robots do the jobs of propagating information and reproducing themselves so efficiently that human survival becomes largely irrelevant to memes. At that point we may well have designed ourselves into oblivion.

The inevitable spread of parasitic and commensal memes undermines the strongest sociobiological claims that human culture is the faithful servant of human reproductive success. A complicated process of feedback between genes and memes is more plausible. Lumsden and Wilson argue that genes are largely in control of memes; they claim that "genetic natural selection operates in such a way as to keep culture on a leash."62 This is an unintentionally apt metaphor. As most dog owners quickly learn, it is sometimes difficult to tell who is dragging whom around.63

If parasitic memes arise, why doesn't the human body evolve to avoid them? In the biological world, organisms do evolve to resist parasites. Hosts that are easily infected by parasites may tend to produce fewer offspring, so over time natural selection favors hosts that develop ways of preventing infection. But natural selection also creates pressures on parasites to increase their abilities to infect and replicate.⁶⁴ The result is a sort of arms race in which parasite and host attempt to develop newer and more effective measures to produce and prevent infection, respectively.⁶⁵ If parasites can evolve faster than their hosts, natural selection enables them to adapt to their hosts' defenses more quickly than the hosts can adapt to create new ones. Parasites that go through many generations in a relatively short period of time (like bacteria or viruses) are more likely to win an arms race because natural selection works faster on them.66

Of course, hosts have other ways of dealing with parasites. They can create incentives for parasites to develop into commensals or mutualists, for example, by developing a more hospitable environment for variants less harmful to the host. They can even modify their own characteristics so that their guests are less harmful to them. And, as we have seen, there are also evolutionary pressures on parasites to rein in their harmful effects. If a parasite is too virulent it will destroy its host too quickly and lessen its chances for future transmission to new hosts.⁶⁷ So pressures to increase infection rates compete with pressures to become less harmful to the host—at least before the parasite has transmitted its genes.

Memes are constantly mutating and recombining; they evolve much faster than human genes can. Thus they would almost always win an arms race with the human body. But there is an important difference between the biological and cultural worlds. Memes are usually transmitted to people already constituted by many existing memes and meme complexes. Indeed, most memetic infection is possible only because human minds are already infested with other memes—for example, linguistic skills. So parasitic memes do not simply invade an unaided human host; they compete against an army of cultural software that can adapt more quickly than human genes.

Thus human beings develop memetic filters to ward off potential cultural parasites. Education, for example, can enable us to discriminate between useful and harmful ideas and to ward off bad influences; we can use our powers of reason to overcome our prejudices and persuade others to do likewise. Our cultural software is a bit like an immune system, which attempts to weed out virulent infections. Sometimes the immune system does not recognize the invader as a danger, sometimes it is overtaxed by the infection, and sometimes it overreacts to a harmless invader, as in the case of allergies.⁶⁸ Like immune systems, our cultural software will never have perfect information—it will always engage in rules of thumb, encouraging infection by those memes most likely to be beneficial and blocking or neutralizing those that might be harmful.

All of these defense mechanisms have an interesting effect: they fundamentally change the nature of the organism being defended. People ward off some memes by incorporating others. In the process, they become cultural beings, interested no longer simply in the reproduction of their genetic information but also in the promulgation and protection of their beliefs, values, and skills. Human beings resist culture only by allowing themselves to be conquered by it.

If our memes do affect our behavior, one of the most important ways is by promoting their own propagation. Space in the minds of human beings is limited. So is the time needed to read the books, listen to the music, and learn the motor skills involved in successful cultural transmission. There is not only extensive competition among memes for space in the minds of prospective hosts, but also competition within any existing host for behavior devoted to propagation.⁶⁹ Thus, some (but not all) of our behavior can be seen as the demands for the reproduction of our various memes, just as some (but not all) of our behavior can be seen as responding to demands for the reproduction of our genes.

Many forms of human behavior seem consistent with this hypothesis. Parents not only want their children to survive; they want to pass along their culture and religious beliefs to their children as well. Parents do not want this merely because they believe that these skills and beliefs will enhance their children's future reproductive success; they also want to perpetuate their own religious and cultural beliefs. People can thus be cultural parents as well as biological parents, and bonds of cultural transmission (as in adoptive or foster parenting) can often be strong indeed. Biological parents whose children are raised by others may feel a sense of loss for many reasons, but surely one is the failure to pass along their values and beliefs.

People transmit cultural software not only through the family but also through education. Teachers and mentors can sometimes have quasi-parental interests in the success of their students and protégés, their intellectual offspring. Academics sometimes fight heatedly about hiring and tenure decisions because they want to ensure that people with similar disciplinary commitments succeed them.

Most important, people often seem to have a deep interest in propagation of cultural beliefs to total strangers, as is demonstrated by religious proselytization. If culture were simply a domesticated pet on Lumsden and Wilson's genetic leash, we would expect that people would pass their most treasured memes only to their relatives, as a sort of "family secret" that would benefit future reproductive success. 70 In fact, we see quite the opposite phenomenon. People are often very interested in the propagation of their cultural software in perfect strangers, whose minds they view as fertile ground for the spread of memes.

This motivation is partly explained by the benefits that come from social coordination. I may have interests, for example, in ensuring that everyone speaks the same language and drives on the same side of the road as I do. But much of our discomfort with cultural differences cannot be explained in this way. Much proselytization cannot simply be seen as a desire to solve collective action problems. There are real advantages to being in the cultural majority, but they stem from the fact that majorities tend to take care of their own common interests, usually to the detriment of cultural minorities.

Altruistic behavior between people who share similar cultural software and oppression or neglect of those with different cultural software would make sense if one goal of human behavior were to propagate memes. Evolutionary biologists argue that competition between genes sometimes leads to cooperative behavior between individuals that maximizes the reproductive success of their commonly held genes.⁷¹ We might expect to see the same thing in the world of culture. Religious groups, academic disciplines, and political parties may help spread and preserve memes more efficiently than individual action.⁷²

The flip side of ideological and religious conflicts, after all, is relative agreement and cooperation within each of the warring sides.

Just as individuals have varying degrees of genetic kinship, they also have varying degrees of memetic kinship.⁷³ The two forms of kinship are crosscutting: people can have many of the same memes even if they are completely unrelated. If the analogy to evolutionary arguments about kin-based altruism holds, then we would predict considerable altruistic behavior between people with lots of similar cultural software—for example, people of the same religion or culture, teachers and students, members of the same fraternity or club. Of course, just as in the case of genetic explanations of altruism, not all altruistic behavior can or should be explained in this way.

We often see people energetically promulgating their memes in the forms of beliefs, behaviors, artifacts, and customs while struggling with others who resist or disagree. Just as competition between biological kin groups can lead to strife, so can competition between cultural kin groups. The history of humanity is littered with religious wars, ideological conflicts, and partisan disputes, many of which cross lines of genetic kinship.⁷⁴ Within the tiny world of the academy, participants jealously guard their turf and promote their own disciplines and approaches, often with a violence seemingly out of proportion to the importance of the struggle. It is often said that such conflicts are so bitter because so little is at stake. From the standpoint of cultural evolution, however, one might say that a great deal is at stake: control over the reproduction of cultural software. If memes are programmed to survive and reproduce, such struggles are serious business, at least for them. Kulturkampfs—or cultural struggles—can be seen quite literally as competition between different meme complexes struggling for mastery and survival, using their human carriers as the means of carrying out this struggle. We can even give a memetic spin to the Gramscian idea of hegemony. Cultural hegemony, we might say, is control over the means of memetic reproduction.

Memes in Conventions and Institutions

Memes are the building blocks of institutions and conventions. As ongoing practices of understanding and behavior, institutions and conventions produce new memes. But more important, they also reproduce the memes necessary to make them ongoing practices of understanding and behavior. Institutions and conventions involve meme-making memes—they coordinate memes to reproduce themselves.

Take, for example, the institution of a club or a lodge.⁷⁵ The institution of a lodge usually involves memes for common practices or rituals that distinguish its members, memes for behaving altruistically toward other members, and

memes for gaining new members to continue the practices and rituals of the lodge. If these memes are properly adjusted to each other and to their environment, the entire complex of memes will be self-perpetuating. People will continue to join the lodge over many years, and its rituals will be perpetuated in its members.

A second example is an annual lecture.⁷⁶ An annual lecture series involves a series of coordinated skills that produce new memes (the lecture) as well as perpetuating the skills necessary to perpetuate the institution. Put another way, one has to know how to put on an annual lecture series, and this coordinated body of cultural know-how by various people constitutes the institution. The selection committee chooses a speaker every year, the treasurer raises funds, the publicity committee issues invitations to the guests, the guests show up and sit in the audience, the chair introduces the speaker, and the speaker prepares a set of remarks and gives the lecture. Through the coordination of these various skills, new memes are distributed (those in the lecture and in the brochures, for example), but more important, expectations are created for the perpetuation of the lecture series in the following year.

In both of these examples, the meme-making institutions and conventions depend on a delicate ecological balance that requires coordination between cultural skills and adaptation to the social environment. If parts of the coordinated understandings and actions fail to occur, or if they misfire, the institution can come grinding to a halt; it will fail to produce new memes, in particular the memes that ensure its reproduction. The lodge may be too picky in its membership requirements, for example, or the treasurer may fail to raise sufficient funds for next year's lecture. If the environment in which the memes perpetuate changes too much, the institution may find itself unable to reproduce. Thus a lecture series may fail to gain an audience because of other forms of entertainment, such as television or movies. Conversely, over time, the memes produced by the institution may change. A lodge that began as a social club may turn into a charitable organization, a lecture series that began as a popular exposition of recent scientific discoveries may become a more serious academic event. Many ancient institutions are able to change their rituals, practices, and beliefs in significant ways and yet retain their self-perpetuating character. Such a task is no small feat, for if the complex of coordinated memes changes too much or too quickly, it may disturb the equilibrium of cooperation that assures its continued reproduction. Yet significant transformations do occur, especially over long periods of time. The older versions of the institution are linked to the present one less by clear resemblance than by a line of memetic descent.

Conventions and institutions can change over time, but they also resist change. An important part of cultural conventions and institutions involves expectations about how other people will behave.⁷⁷ These expectations not only help coordinate behavior, they also stabilize and police it. Some degree of creativity and variance is always possible within conventions and institutions, but too great a variance defeats expectations, threatens stability, and endangers cultural reproduction. Hence great change produces resistance. Of course, strong resistance to interferences with successful reproduction is precisely what we would expect from self-reproducing entities that emerge from natural selection.

Some rational-choice theorists have tried to explain social conventions and institutions as coordinated behaviors of rational actors that are able to surmount collective-action problems.⁷⁸ An example of such a problem is a prisoner's dilemma, in which fear of loss from defection by others tempts parties to defect from coordinated action that might benefit them all. Rational-choice theorists have tried to show how such coordinated behaviors might arise spontaneously. The theory of cultural software approaches this problem from a slightly different perspective. Instead of focusing on how rational behavior of human beings might overcome collective-action problems, it focuses on the collective-action problem faced by memes themselves. We can make an analogy once again to genes that cooperate with each other to create multicelled organisms, thereby enhancing their joint survival chances. Conventions and institutions are coordinated complexes of meme-making memes. They cooperate with each other because this cooperation enhances their joint reproductive success. By assisting in each other's reproduction, each meme helps the other gain precious space in human minds and bodily behaviors.

Conventions and institutions reproduce expectations in people about how others will behave. These expectations are essential parts of self-reproduction. They are important because they let conventions and institutions "turn the tables" on the problem of collective action. Once cultural expectations are in place and continually reproduced in human minds, it takes collective effort for variant memes to overcome the settled body of cultural expectations.

Why switch our focus from the rational actor to the unit of cultural transmission? I noted earlier that memes can be mutualist, commensal, or parasitic with respect to their human hosts. Because commensal and even parasitic memes are an almost inevitable development in cultural evolution, we cannot assume that all conventions and institutions are merely solutions to coordination problems that benefit human hosts, either in terms of their reproductive success or their interests as rational actors. Some conventions and institutions may be commensals, and some may actually be quite harmful.

The history of human culture is the history of human susceptibility to various kinds of memes. As people are infected, their memes prepare the way for new memes, that, in the process, alter and even increase human susceptibility to memes. Consider, for example, the types of memetic infections made possible by learning a human language. These new susceptibilities are passed along from generation to generation in the form of human culture. At some point in human history, human beings became susceptible to a variety of conventions and institutions. They became fertile ground for any number of selfreproducing complexes of memes. This development may originally have been a good one from the standpoint of human reproductive success. But it made human beings susceptible to infection from many different kinds of conventions that did not necessarily have the best interests of humanity at heart. Of course, internalized memes do not merely weaken human immunity to new invasions of memes. They also create a new "immune system" that can ward off some harmful memes. However, this cultural immune system cannot perfectly distinguish between memes that are useful in the long run and those that are not.

Social conventions and institutions are possible because our brains developed so that they were fertile soil for certain types of self-perpetuating skills. But once this fertile ground was created, it became suitable for many different kinds of conventions that might be created in the future, just as fertile soil can admit weeds as well as useful plants. So we cannot assume that all conventions are beneficial to the members who engage in them. Some conventions (for example, slavery) are indeed "solutions" to social coordination problems, but they are not necessarily beneficial solutions. Other conventions, like the meanings of certain fashion designs, are commensal, in that they have very little benefit.

Shifting our attention from the interests of rational actors to the "interests" of conventions and institutions themselves—as collections of self-reproducing memes—puts a very different spin on the growth and development of human culture. It allows us to understand evolutionary developments without having to explain them in functionalist terms or in terms of rational benefit to humanity. We need no longer conjure up "just-so" stories to explain all of the various features of human culture. Instead, we can understand human culture as a compromise and conflict between the interests of persons, their genes, and their memes. We can make better sense of the idea of conventions or institutions that literally take on a life of their own, regardless of their current or long-term benefit to humanity.

This evolutionary approach to conventions has a further advantage. Dawkins's original formulation of memes was an extension of his theory of "selfish genes," which argued that genes used organisms to maximize their own reproductive success. Naturalists like Stephen Jay Gould and Richard Lewontin have countered that too exclusive a focus on the gene fails to reckon with the constraining force of the architecture of organisms.⁷⁹ Gould and Lewontin's attack on the "adaptationist program" reminds us that "organisms are not so much paragons of design as compromises of design."80 Natural selection cannot always perfectly hone organisms to maximal reproductive success in their ecological niches. Organisms may lack the necessary genetic variation for natural selection to do its work. Previous design choices may limit future innovations, allowing only bricolage rather than bottom-up design.

A similar point applies to the complexes of memes we find in conventions and institutions: they too are compromises of design. Only some changes are possible if conventions and institutions are to maintain their self-reproducing character. And, like species, they also face "architectural" constraints. They can only change in certain ways, given their previous history and the cultural means at hand. Hence, conventions and institutions produced through cultural evolution are highly unlikely to perfectly optimize any external standards of design in their current environment, whether that standard be social functionality, moral efficacy, or economic efficiency.⁸¹ Indeed, the imperfections and the jerry-built character of conventions and institutions are the best evidence of their historical development. This realization connects the theory of memetic evolution with my discussion of cultural bricolage in Chapter 2.

We cannot always infer the current utility of a feature of a convention or institution from its current existence. And we cannot infer from a feature's current utility the reasons for its origin. Rather, we are likely to see, in Gould's phrase, "panda's thumbs" in both conventions and institutions. Cultural developments of social conventions and institutions are likely to be extapations, in which memes adapt old features to new uses in a changed environment. Features of existing conventions and institutions may often have arisen for one reason, but now serve very different functions and purposes.

The argument I have just presented synthesizes Dawkins's concept of memes with Gould's theory of architectural constraint. The concept of memes was originally coined by Dawkins, whose views about adaptation have been criticized by Gould. But once we take into account the role of meme complexes in cultural evolution, and the need for these complexes to reproduce together in a given ecology, we see that Gould's point about evolutionary bricolage applies with equal force to models of cultural evolution. And this requires us to modify Dawkins's original conception of memetic evolution.

Indeed, there is something entirely fitting in bringing these two seemingly disparate strands of evolutionary theory together. The original inspiration for Gould and Lewontin's theory of evolutionary bricolage was an example drawn from the world of culture—the spandrels in the Basilica of San Marco. Spandrels are triangular spaces created when one places a cathedral dome on top of four rounded archways set at right angles to each other. Artists made use of these spaces to place elaborate mosaics and other decorations. Gould and Lewontin pointed out that it would be fallacious to assume that the basilica was

designed to produce the decorative spandrels; rather the custom of decorating the spandrels is simply an ingenious use of spaces that necessarily resulted from previous limitations on the design of cathedrals.82 In the same way, if we wish to study the development of cultural software, we must not assume that all features of human thought and practice are currently or perfectly adaptive to memetic survival. Rather, we must patiently investigate the ways in which layers of memetic innovation occur given the existing constraints of human thought and human cultural conventions.

Because memes can be commensal or even parasitic, we no longer have to explain the development of culture in terms of what is functional for human beings or even for society as a whole. We can shift our focus from what kinds of memes would help human beings or cultures survive to what kinds of memes are most likely to survive and propagate in human beings and their information-processing technology. In many cases, memes are successful replicators because they are true beliefs or because they provide skills useful to human beings. But they need not point toward truth or possess great utility in order to survive and propagate in human minds. They may just as easily spread by playing upon our worst instincts, by pandering to our coarsest or basest desires, by permitting us to avoid recognizing our moral responsibilities, or by encouraging sloth, avarice, and a hundred other vices. Finally, they may proliferate without bringing either significant good or evil into the world. They may multiply simply because they are entertaining or diverting.

Why Memes Survive and Spread

What makes some memes more successful in their environment than others? We can identify three basic kinds of factors. The first are substantive factors involving content. Second are psychological factors—the cognitive structure of human minds and their comparative susceptibilities. A third set of factors is ecological—they concern the nature of social institutions, methods of storing information, and technologies of communication. These different features are linked in practice. For example, the kinds of substantive content that make memes more attractive or more often discussed may depend on

structural features of the human mind and existing religious or educational institutions.

Substantive Factors

Memes tend to spread if they are salient, relevant to existing activities, attractive, or entertaining, or if they generate strong emotions. Sometimes it is not difficult to see why some memes spread more than others. Jokes and skills like juggling or playing a musical instrument are widely distributed because they are entertaining. Other memes spread because they are relevant to many different people's lives and interests. Consider as an example the number of songs about the various aspects of love and courtship.

Memes improve their reproductive success if they have behavioral effects that promulgate their own spread. A catchy melody, for example, may cause people to hum or sing it repeatedly, thus increasing the number of times that it is heard by other people. A good joke spreads rapidly because people enjoy telling it to others. Memes may be more successful if they encourage proselytization, appear to be beneficial (thus encouraging sharing with friends and relations), promote cooperation with others, and hide any maladaptive features for as long as necessary to spread widely.1

Another strategy for survival is to disable or preempt potential competitors in the environment.² Standard examples are memes for faith, which discourage skeptical beliefs and the sort of critical inquiry that would tend to dislodge faith.3 Ideas of tolerance or free expression tend to assist their own propagation, but they also assist many other competing ideas as well, including ideas of intolerance and censorship. Complexes of memes working together may create joint defense mechanisms. Examples are warnings in chain letters that if recipients break the chain something bad will happen to them, and rumors of powerful conspiracies that explain objections on the grounds that all objectors are either part of or have been hoodwinked by the conspiracy.⁴ To this one might add theories of ideology that make use of concepts of false consciousness to dismiss critics.

Still another method for memes to improve their chances of reproductive success is to attach themselves to other successful memes. Religions, for example, usually include many accretions over time. These accretions benefit from the general acceptance of religious belief and powerful memes for faith. Believers follow the tenets and practices of a general religious tradition together rather than investigating each one separately. Of course, the meme that lives by linkage can also die by linkage. If a meme is too closely linked to others that lose favor, it may be filtered out precisely because of these associations—a memetic baby thrown out with the bathwater.5

One might think that the most important factor in increasing a meme's reproductive success would be its truth or falsity. But many memes cannot be said to be true or false. Examples are bodily skills like dancing, and practical skills like those involved in being polite. Informational and cognitive filters, which shape thought, are among the major concerns of the theory of ideology; they are neither true nor false, though they can produce true and false beliefs. Finally, many philosophers hold (incorrectly, in my view) that statements of political, moral, and aesthetic value cannot be true or false. For these noncognitivists, truth or falsity is irrelevant to the success of a large number of memes.

Even with respect to memes which correspond to beliefs that can be true or false, there are several reasons why truth does not necessarily increase reproductive success and falsity does not necessarily diminish it. First, when a belief is obviously true, no one may pay much attention to it or think about it. As a result, it may be much less likely to be communicated to others. Memes so obvious that they are rarely discussed tend to lie dormant in minds; they are, quite literally, things that go without saying. Regular and prolific replication often matters more to reproductive success than durable presence in a particular human mind, because a particular carrier might die or forget. Hence memes may be more successful if they are controversial, taking that word in its literal sense as that which produces conversation.

Second, not all of the true things we believe are actually recorded in our minds at any point. For example, most people probably believe that there are no indigenous palm trees in Antarctica, but it is likely that they have never thought about it before the fact was brought to their attention. Many things we "believe" in the ordinary sense of that word are inferable from other beliefs that are stored mentally. Because many if not most of our true beliefs are of this form, it may be quite important for true beliefs to be generated, used, and thought about if they are ever to be spread to others. Even true beliefs that many people could generate independently will not be generated and spread unless occasions arise to generate and spread them.

Third, true beliefs are much more likely to be communicated in response to false beliefs or only partially true beliefs (approximations of the truth, for example). This suggests that some false and true beliefs are coadapted: the presence of one spurs the communication and spread of the other. There is an analogous problem for religious beliefs. Religious faith can weaken over time if it is not occasionally faced with challenges. Hence heresies and external opposition to faith may sometimes increase religious fervor, proselytization, and the propagation of religious memes.⁷

Fourth, some memes may be employed, generated, or communicated precisely because it is difficult to tell whether they are true or false. Many of the most commonly communicated ideas are those whose truth and falsity cannot

be determined, which is why they are the subject of endless debate. This debate, in turn, ensures their continual transmission and survival. An analogous point applies to questions of practical reasoning and aesthetic judgment. A course of action is most likely to be debated precisely when its consequences and appropriateness are unclear; some works of art improve their chances of success by being controversial.

Fifth, beliefs that are clearly true often have unequivocal meanings or unequivocal applications; otherwise they would not be clearly true. But such clarity may not improve their reproductive success. Some memes are more likely to reproduce themselves if they are ambiguous—if they mean different things to different people, or even to the same person. This is especially true in the world of values. Principles like equality and liberty are ambiguous in their reference and hence can be—and are—invoked by different sides of a dispute. They become objects of struggle, and through this struggle they are repeatedly communicated and transmitted, thus ensuring their continued survival. In like fashion, the most heavily litigated and discussed parts of a legal code or constitution are often those that are least clear, or that become increasingly unclear through successive judicial interpretations.

Finally, truth or falsity may not be relevant to survival because we can remember and transmit beliefs even if they are false, bigoted, or unjust. Some beliefs survive precisely because they are understood to be false or wrong. They are helpful examples of falsehood or wrongfulness that are continually repeated because of their helpfulness.

Psychological Factors

Many of the most important factors in the spread of memes depend less upon their substance than upon features of the human mind. We have already noted one such factor—the capacity of a symbol or belief to raise strong emotions. Memes better adapted to the architecture of the mind take root more readily than others; hence we can study their comparative success for clues to the nature of this architecture. Experiments have shown, for example, that human beings develop certain basic level categories like "bird," which are easier to remember and employ than more abstract concepts like "flying thing" and more concrete concepts like "yellow-bellied thrush." These basic level categories are more "catching"; studying rates of comparative "infection" gives us important clues about the organization of the mind.⁸

MEMORY AND COMPREHENSION. Human memory storage is an inevitable bottleneck for cultural transmission. Hence one of the most important factors affecting the survival of memorization. Ease of memorization. Ease of memorization.

orization depends on complexity, but complexity is not an inherent feature of information. It is partly a function of mental architecture. Human minds are not general-purpose memorization machines. They have particular strengths and weaknesses that are the result of prior evolutionary pressures and compromises of design. Different kinds of memes and complexes of memes face different degrees of success in this architecture.

Compare the memorizing abilities of a computer with those of a human being. What is easy for a human being to remember may be difficult for a computer, and vice versa. Computers can easily memorize long stings of numbers that would tax any human memory. Narratives and myths are effective methods for human memorization, but not necessarily effective methods of computer memorization. Human beings can easily store hundreds of tales and myths that can be told in multiple variations. It is much easier for human beings to remember and recite a story than to remember and recite a text of a story word for word. On the other hand, it is very difficult to store a myth on a computer, although we can easily provide it with different textual versions of a myth. Tales and myths are well-designed vehicles for human memory storage; this explains why they remain useful aids to memory to this day, and why they have survived without being forgotten. It is even possible that there were evolutionary advantages for human beings to storing information in narratives. The memorizability of narratives suggests both the internal structure of human memory and the important ways that it differs from those of currently existing computers.

Ease of memorization is especially important in oral cultures that have not developed writing or widespread literacy. In oral cultures, information that cannot be put in easily remembered forms will likely be forgotten. Hence the importance of bards and storytellers, who serve as walking encyclopedias. In oral cultures, songs and stories do multiple duty as popular entertainment, literature, history, religious doctrine, and canons of social instruction. As a result, branches of art and learning are not strongly differentiated. 10 Successful memes must attach themselves to easily remembered forms like stories, songs, and bodily movements, just as medical students to this day learn complex anatomical lists through the use of acronyms. Memes that are hard to remember either will be forgotten or must be transformed into more easily remembered forms before they can be widely spread throughout a culture.¹¹

The invention of writing revolutionizes the cultural environment. Human memory is less of a bottleneck for memetic survival, because it can be supplemented by external memory storage. New forms of literature can develop and may even supplant those found in the oral tradition. Put more generally, every new communication technology leads to new and different susceptibilities for memetic infection; it creates a new ecology for memetic growth and reproduction. Changes in the ecology mean that rates of differential survival and reproduction change; new memes develop that could not have survived or reproduced as plentifully in the earlier environment.

This insight allows us to connect the theory of cultural software with the theory of media analysis. Media analysts like Marshall McLuhan and his followers argued that changes in dominant forms of communication (and hence memory storage) lead to changes in human thought and human culture. Put in terms of the theory of cultural software, changes in media are changes in ecology; they create new selection pressures for memes that lead to new and different kinds of cultural software in human minds. In particular, the movement from an oral to a written culture, and then to a televisual culture, has had significant effects on human memory and hence on human thought and culture.¹² Media analysts argue that styles of thought and expression differ markedly in oral and written cultures. Oral cultures feature thought that is figural, repetitive, concrete, and diffuse; in written cultures, thought tends to become more conceptual, linear, abstract, and analytic.13 The latter kind of thought emerges precisely because print media permit it. In like fashion, forms of thought and expression start to change again as television begins to dominate communication.14

The subsequent move to a computer-oriented information society will doubtless further change our ability to store and process information, again revolutionizing our culture and our forms of thought. We are already seeing the signs of this in the information explosion that accompanies computerization. This explosion not only increases the life chances of many different kinds of memes; it also creates the need and the opportunity for ever new forms of filtering to control the amount of information being created and broadcast. As a result, in the information age, filters increasingly determine what information we receive and how we receive it. In the age of information, the filter is king.

The details of representation are sometimes as much a candidate for natural selection as the context represented.¹⁵ Memes become coadapted to other memes that help in their delivery and memorization, just as information had to be conveyed in narrative or poetic form in oral cultures to ensure memorization and comprehension. In the relentless competition for human memory space, certain methods of communication win out over others: messages coded in rhymes or pithy sayings are memorized better than other messages; commercials with flashy graphics and news reports that resemble entertainment programs garner more attention than less entertaining forms. Media critics have documented how television has tended to merge news, political coverage, and entertainment, and a similar process appears to be happening in media coverage of the legal system.16

Related to ease of memorization is ease of comprehension. Human beings

are less susceptible to memes that they do not understand. Different minds have different degrees of susceptibility to memetic invasion, depending in part on their education and experience. A text that is easy for someone already trained in a discipline may be difficult for a lay person; a sentence easy for a native speaker to comprehend may be more difficult for another person. People who are immune to written language may nevertheless be susceptible to memes expressed in television shows, movies, or music. Different rates of comprehensibility create selection pressures on memes to be expressed in easily communicated and digested forms. Otherwise, memes must content themselves with smaller ecological niches—for example, in subcultures like academic writing.

Like memorizability, comprehensibility is often greatly affected by the medium of communication. Print media make much greater demands on comprehension and require more sustained attention than television. Television has a further advantage: it makes information entertaining by using music, quick image changes, and flashy graphics. The different features of these media have two separate types of effects. First, they bestow a survival advantage to memes conveyed on television, although there are compensating disadvantages as well-for example, televisual information may be viewed as disposable and hence more easily forgotten. Second, because television can be entertaining and absorbing in ways that print media cannot, there is continual selection pressure in television for memes to be more and more entertaining and absorbing. More entertaining programming tends to weed out less entertaining programming. Certain types of broadcasts—for example, a stationary camera focusing on an extended lecture by a standing speaker—tend to be weeded out because they are not "good television." More generally, memes involved in public discourse tend to become coadapted with memes that are optimal for communication on television, producing important alterations to both.

On television, certain styles of communication tend to dominate others: For example, in the current environment, at least, ten-second "sound bites" seem better adapted to the demands of television than four-hour discussions of policy issues. ¹⁷ Ideas embodied in pictures and accompanied by music tend to dominate ideas conveyed through rolling black text on a white screen. More generally, memes conveyed through a medium's favored forms of communication tend to thrive; memes that cannot be as effectively conveyed in this fashion tend increasingly to disappear from television broadcasts. This competition affects content as well as form. Political discourse has long since begun to borrow heavily from advertising; politicians have learned to stage media events that grab precious television time. Because television favors entertainment, there are selection pressures on public discourse, advertising, and even coverage of the legal system to conform to these standards and increasingly to

resemble other forms of television entertainment.¹⁸ All of these tendencies confirm the role of natural selection in the development of culture.

EASE OF COMMUNICATION. Memes that are easy to communicate tend to spread more than those that are more difficult to communicate. Ease of communication is not necessarily the same as ease of memorization or even ease of comprehension. A list of numbers may be easy to communicate but difficult to remember. A deeply personal experience may be easy to remember but difficult to communicate.¹⁹

Every teacher knows that some ideas are more difficult to convey than others. Listeners often take away misunderstandings of complicated ideas because the misunderstandings are easier to comprehend and communicate to others than the original, more complicated idea. As a result, the distorted or mutated version may spread more widely than the original. Indeed, repeated communication can affect not only the substance of communication but its form as well. Some words are harder to pronounce than others, leading to mutations of pronunciation.²⁰

Unlike genetic transmission, which engenders relatively faithful copying, cultural transmission normally involves alteration and mutation. Hence in explaining the spread of shared cultural software, we must account both for the ability of cultural software to spread and its ability to preserve some measure of identity.²¹ Because opportunities for alteration are so commonplace, the most widely shared features of our cultural software are those that can best resist alteration after repeated transmission and mutation.²²

Stories provide a good example. Each time a story is told, it is likely that the version is slightly different from the last. Some details may be added, others subtracted, and still others compressed or merged. Only the most easily communicated, understood, and remembered features tend to be preserved.²³ Most people who remember the biblical story of Joseph, for example, believe that Joseph was sold into slavery by his brothers. In fact, the story told in the Bible is more complicated and thus less easy to remember. The Hebrew text suggests that Joseph's brothers threw him into a pit. He was rescued by some Midianite merchants, and they sold him into slavery. But the "folk" version of the story has become more widely transmitted than the original.²⁴

In similar fashion, statements and slogans tend to be transformed through repetition until they are relatively easy to remember and transmit to others. This may help to explain the familiar phenomenon of famous "quotations" that were never actually spoken but are variants of what was actually said.²⁵ Not surprisingly, political slogans spread more easily than the complicated political theories that they stand for, and they have the further survival advantage that they stand for many different things to many different people.

This evolutionary account explains why a wide variety of cultures have similar narratives and myths. Claude Lévi-Strauss argued that myths in different cultures were transformations of basic narrative structures that in turn reflected basic structures of the human unconscious.²⁶ But we can explain matters more simply. The "universal structures" that we see in human myths and legends may reflect those elements of stories that best survive the continual mutation and alteration that comes with repeated tellings. Moreover, because the spread of myths depends on the ecology of human minds, their content and structure may shift over time.

REFLEXIVE BELIEFS. One of the most important factors in human susceptibility to memes is the reflexive nature of our thought. People not only can have ideas, they also can have ideas about ideas. They can have attitudes or opinions about particular beliefs and ways of thinking. For example, people can understand ideas without being convinced of them; they can believe that certain things are not true; they can recognize that certain opinions are odious. They can engage in mental simulations, plan, exercise foresight, imagine, model, play, or fantasize.²⁷ People often remember memes precisely because they are false, wicked, or don't work. Parents take great pains to teach children what not to think and what not to do, and, if they are lucky, their children internalize these lessons.

People not only can produce and store interpretations of events, they can produce and store interpretations of those interpretations.²⁸ For example, historians not only develop interpretations of the American Revolution, they also remember and discuss the various interpretations of other historians about the Revolution. Moreover, they can pass these interpretations on to their students and other historians even if they don't accept them.

The recursiveness of human thought makes people susceptible to many more types of memes than they actually accept, believe, or act upon. Memes may not die out even if people reject or disbelieve them, because people can still remember and discuss them—with the admonition that "this is wrong" or "this doesn't work."²⁹ False ideas and bad practices can remain in human memory even though they are known to be bad or false. What is stored in memory can be communicated to others. As a result, false ideas and harmful cultural skills can be passed on to new generations despite their being known to be false or harmful. These memes can live to another day, when they can significantly affect the behavior of another host. In such ways, superstitions and prejudices can survive even though people decisively reject them. A more benign example takes the form of historical interpretations rejected by one generation of historians that are retained in historiography and eventually regain favor in a subsequent generation.

Finally, people can store ideas and beliefs even if they do not completely understand them and are not certain whether they are true.³⁰ People may believe that space curves near a heavy mass, for example, because they read it in a book, although they really don't understand how this could be so. They can retain such beliefs pending further information that might clarify the beliefs or demonstrate the beliefs to be true. And people can hold these beliefs indefinitely, even if no additional clarification or proof is forthcoming.³¹

Such half-understood beliefs are not restricted to obscure scientific theories. People can hold beliefs about UFOs or religious doctrines, for example, whether or not they fully understand or know the truth about such things. In fact, people may be particularly susceptible to what is mysterious precisely because mysteries resist solution or comprehension.³² Exposure to ideas that are difficult to prove or comprehend may even encourage their being discussed or talked about further. Their very inconsistency with other beliefs and their very inability to be fully comprehended make mysteries intriguing and attractive and lead to their further discussion and distribution into other minds. In this way, an otherwise beneficial feature of human cognition—the ability to store and reconsider incompletely understood information—creates the opportunity for the differential reproductive success of a certain kind of meme—the mysterious—in the environment of human minds.

Ecological Factors

In most cultures, the reproductive success of memes is largely determined by other memes and by the institutions that use and propagate other memes. Previously internalized memes shape mental susceptibilities to new memes; the cultural skills involved in institutions create the environments in which memes compete. Thus the pool of existing memes creates the basic ecology for other memes. Cultures are like the tropics, where the landscape is overgrown by plant and animal species, and where chances of survival and reproduction are largely determined by the ecology of other organisms rather than by the original physical habitat.³³ Tropical climes are well known for their intricate ecosystems and for the strange and freakish creatures that they produce.

In short, we should think of cultures as ecologies rather than as well-integrated and organic unities. They are inherently open systems rather than closed ones. Cultures involve an ecological equilibrium between different forms of cultural software, an equilibrium that may be disturbed, reconfigured, or even destroyed by memetic invasion or environmental disturbance.

SEXUAL SELECTION AND BANDWAGON EFFECTS. The crush of animal and plant life in diverse ecologies creates opportunities for exaggerated and bizarre

traits. This is due in part to an evolutionary phenomenon called sexual selection. In the natural world, females tend to choose mates based on characteristics that are attractive to other females. They do this to guarantee that their male offspring will be equally attractive to future generations of females, for offspring that attract no mates will produce no offspring of their own.

Females look for characteristics in mates that tend to correlate with the reproductive success of their offspring. Once female preference for a feature is generally established, however, the feature by itself makes the offspring more desirable to future females. Thus females want the feature in their mates simply because all other females also want the feature. The result is a "runaway" effect: the preference for the feature is intensified out of proportion to its otherwise beneficial effects. Thus female peacocks prefer peacocks with long bright tail feathers. These features may confer no present additional evolutionary advantage—they may even be debilitating to the male—but because of sexual selection they increase the chance that these males and the females who select them will reproduce their genes in future generations.

In the world of culture, analogies to sexual selection can occur in several different ways. First, to some extent, individuals can choose what beliefs and cultural skills they will internalize. They may choose to adopt beliefs and behaviors of powerful and influential people because they believe that this selection will make them seem influential and powerful. This process can snowball so that status-seeking individuals attempt to outdo each other in cultural displays. The result is extremism in belief and behavior, because the extremist, like the long-tailed peacock, seems to be at the leading edge of a trend. The desire to be thought highly successful, powerful, or pious can even lead to competitive construction of elaborate cultural monuments, like pyramids and cathedrals.

Second, a cultural equivalent of sexual selection produces "bandwagon effects." People may engage in faddish beliefs or behaviors because they believe that others regard them as desirable, and the belief that others find them desirable increases their desirability even more. John Maynard Keynes's famous description of the stock market as a beauty contest is based on a similar logic—people often buy stocks because they believe that others value them, and this drives up their value out of proportion to a company's expected future earnings. Signals and filters can play important roles in producing bandwagon effects. Best-seller lists are institutional filters that use people's past buying decisions, but they simultaneously act as an advertising gimmick by signaling other people's preferences. Once a book sells enough copies to get on the best-seller lists, its sales may increase rapidly.³⁷

Third, sexual selection can occur in the way memes form alliances with

other memes. Just as females seek to mate with males whose offspring will be desirable to future females, memes may face evolutionary pressures to join forces with memes that seem particularly successful in gaining entry to human minds. These traits can also snowball. Suppose that flashy graphics, loud volume, and quick cross-cutting of images tend to attract the attention of television viewers. Then memes may come to be delivered through increasingly flashier graphics, louder volumes, and quicker cross-cutting. This may explain the evolution of some forms of television advertising.

INSTITUTIONAL AUTHORITY. An important feature of human culture is that human beings can accept beliefs and adopt customs and practices because of institutional authority. We believe many things not because we have direct evidence for them or have gone about proving them to our own satisfaction but because they have been communicated to us by people and institutions we trust. Similarly, there are many practices and customs that we have adopted not because we have independently determined that they are optimal but because other people engage in them. Sometimes doing what others do has independent advantages—for example, cooperation and coordination can sometimes solve collective-action problems to the benefit of all parties. However, not all examples of following what others do can be explained or justified in this way. Driving the same car that everyone else drives, following the current fashion trends, or hewing to the party line does not necessarily solve collective-action problems.

Following the dictates of institutional authority makes sense for a different reason. Many things cannot be demonstrated for certain, and it is often difficult to know what course of action is best. Hence it may be rational for people to believe things simply because that is what other people believe, and to do things simply because others do them. Believing and doing these things is rational, not by virtue of their content but by virtue of their source.³⁸ If this is indeed rational behavior, we would expect that people in different parts of the world would have different beliefs and customs because they trusted and learned from different sources of belief and action—the people who educated them.

Cultural traditions have a kind of institutional authority, and a similar logic applies to them. Traditions provide people with things to believe and ways to behave. Traditions are not necessarily antithetical to rational action: people rationally strategize within the norms of their tradition and its beliefs; they can even decide to forsake their traditions for other beliefs and practices. That is one way that traditions evolve. But it may be reasonable for people to hew to traditional beliefs and practices when it is difficult and costly to discover what to otherwise believe or do. This is especially true of problems of practical

reason. The long-run usefulness of practices may be difficult to determine in advance. Hence following tradition becomes a useful means for solving problems of ordinary living.³⁹

INSTITUTIONAL CONTEXT. Memes are more likely to spread if they are relevant to existing institutions, either because they are associated with the institution or because they give rise to appropriate action in the institution. Handel's Messiah, Tchaikovsky's Nutcracker, Frank Capra's It's a Wonderful Life, and Dickens's A Christmas Carol are among the best known and most often repeated of their works because these works are associated with the institution of Christmas. Among the most frequent phrases spoken in many cultures are greetings, comments on the weather, and requests after health.

If an institution requires regular and repeated replication or demonstration of a meme or cultural skill, the chances for survival of that cultural skill are greatly enhanced because the skill is more likely to be remembered and reproduced. Cultural software benefits in particular if there are institutions specifically devoted to its spread and propagation. Examples are schools, churches, libraries, universities, and the family. Some cultures institutionalize the telling of myths and legends, and this helps to ensure their continued survival.

Complicated scientific information depends heavily on institutional structures for its survival and spread. Scientific truths may be quite compelling once demonstrated to an audience prepared to receive them, but they are often difficult to comprehend without considerable training. Hence even the most indubitable of truths may require elaborate institutions of education (including elaborate structures of intellectual authority) if they are to be preserved and propagated. If these institutions fall apart, the true beliefs that they propagate may become extinct as well. Our romantic notion that the truth will out neglects the importance of institutional ecology. Here is yet another example that shows that the truth of a belief does not guarantee its widespread reproductive success; it must find a niche in a suitable environment if it is to survive.

Political beliefs also depend heavily on institutional context, but for somewhat different reasons. Dan Sperber gives the example of the belief that all people are created equal.⁴⁰ This belief is both salient and controversial in societies organized around pervasive social, economic, and political inequalities. That is because the belief has many different implications for such a society. People who like these implications have grounds to accept the belief and incentives to spread it, even in the face of considerable opposition.

This is the memetic version of a familiar theory of ideology—interest-driven explanation. People believe things that jibe with their social, economic, or political interests. The memetic claim is that the institutional environment makes certain people's minds fertile ground for certain types of memes. As a result,

these memes tend to propagate once they are introduced. But if opposition to the implications of a belief is too great, the meme may not spread; at best it may be confined to certain subcultures where it can survive and reproduce.

The memetic account adds a new twist to this familiar explanation of ideology. Because we can model the prevalence of a belief as the result of a competitive equilibrium, the insights of catastrophe theory apply. A slight change in the institutional ecology may have enormous effects completely out of proportion to the degree of ecological change. The belief may spread quickly and unexpectedly. At one point, for example, a particular meme—say one associated with radical egalitarianism—may be able to maintain only a marginal existence in a particular subculture. Yet a slight change in the institutional ecology may lead to an explosive spread of belief. In the new environment, the meme takes off and reaches epidemic proportions.

Nevertheless, if memes are to reproduce widely over time they must be able to adapt themselves to political, social, and economic changes. Thus a meme like equality is most likely to thrive if it can be articulated and adopted by people of different political views over time. Thus successful memes often are subject to wide variation in the form of contrary interpretations and subtle shifts in meaning.

Ideas often change their political valence as they are repeated in new contexts and situations. A good example is the idea that democratic governments should be "colorblind." This idea was associated with a very progressive view of race relations in 1896. It was the basis of Justice Harlan's famous dissent in Plessy v. Ferguson, when he opposed the segregation of railroad facilities. 41 In the 1960s Martin Luther King fought segregation by arguing for an America where citizens would "not be judged by the color of their skin but by content of their character."42 Yet by 1996 colorblindness was the rallying cry of conservatives opposed to affirmative action. A second example involves the libertarian concept of freedom of speech. In the first half of the twentieth century freedom of speech was defended by the political left as a means of protecting political dissenters, minority groups, and labor unions. By the close of the twentieth century it was also being used to defend the rights of cigarette manufacturers, the Ku Klux Klan, sexually harassing employers, multinational media conglomerates, and political action committees opposed to campaign finance reform.43

I call these changes in political valence *ideological drift*. ⁴⁴ They are a ubiquitous phenomenon in social and political life. From a memetic standpoint ideological drift is an example not of political opportunism but of memetic opportunism. As political and social contexts change, slight mutations can make memes newly hospitable to persons who previously would have shunned them. Some members of the American left, for example, have become increasingly

attracted to regulation of campaign spending, pornography, racist speech, and commercial advertising, while conservatives have become increasingly libertarian on the same questions. Shifts in political and social context—as well as in the interests and other beliefs of liberals and conservatives—change the ecology in which political ideas about freedom of speech can thrive. As a result memes may find new minds increasingly hospitable and older hosts increasingly less so. It is important to recognize that memes do not particularly care who invokes them as long as they are regularly invoked. Memes that were once happily nestled in liberal heads will readily and opportunistically mutate to become acceptable to more conservative minds should this increase their chances of propagation and survival.

Shared Understandings and Lines of Memetic Descent

The theory of cultural software holds that individuals share cultural understandings because they possess similar memes. One reason people have similar memes is that they communicated them to each other, or that they live in the same culture and therefore have been exposed to the same memes communicated by other members. But this does not explain how individuals in widely divergent cultures might possess similar tools of understanding, because not all cultures are in continuous contact with each other.

Sometimes individuals have similar cultural software not because they or their cultures have had any recent communicative contact with each other but because their cultural software is descended from a common source. Biological evolution offers a useful analogy. Generally speaking, mammals have four legs and a single head. This common morphology is not the result of crossbreeding between different species but rather is due to the common ancestry of all mammals. The basic pattern for bodily development is passed on in each species even as it evolves and is differentiated among species. That is because biological bricolage is generally conservative, retaining past design choices as the platform for future innovation.

In a similar fashion, the cultural software of present-day human beings builds on the work of previous generations. It is conservative in the same way that biological bricolage is conservative. Earlier forms are retained in later developments, and hence we see many similarities among diverse individuals and cultures to the extent that their cultural software has a common ancestry. Language provides a simple example: Similarities in words across different languages (*father* in English, *Vater* in German, *père* in French, *padre* in Spanish) are evidence of common memetic descent.

One reason for the conservatism of biological development is architectural constraint produced by previous evolution. Previous design choices (like those

in the panda's paw) constrain future morphological development. In the previous chapter I argued that a similar architectural constraint may be at work in meme complexes. Coordinated complexes of memes (like those in a religion) may be able to accommodate only certain kinds of changes if they are to reproduce successfully together. If cultural software spreads and develops through such meme complexes, we might also expect that certain features will be deeply embedded in our cultural software and more resistant to change, just as we would not expect an easy transformation from mammals with four legs and one head to mammals with eight legs and multiple heads.

Furthermore, because evolutionary bricolage must innovate on the basis of existing materials, it tends to retain these materials and adapt and alter them for new purposes. Thus certain tropes, metaphors, symbols, heuristics, or other tools of thinking may run very deep in our culture precisely because they appear so early on in the course of historical development, and therefore have been repeatedly used to fashion later tools through which we presently understand the social world. This depth is not the depth of a core versus a periphery but one produced by repetition and recursion. We can see an instance of this in our earlier etymological example of the word *articulus*, or joint. This word and the concept it represents are used repeatedly to form new words and concepts, which are in turn used to create still other words and concepts, and so on. This process proliferates the original metaphor of joining and dividing into a multitude of later conceptual tools; each of these tools, in turn, is proliferated into new tools, so that the metaphor of joining and dividing appears repeatedly in widely divergent aspects of our cultural software.⁴⁵

On the other hand, it is also possible that certain memes appear in widely divergent cultures not because of a line of common memetic descent but because these cultures faced similar problems and produced similar solutions. For example, Robert Ellickson reports than many different cultures have produced forms of private ownership in land. It is possible that this idea began with a single culture and spread to others because it was useful. But it is also possible that it developed independently in many cultures because people in each culture recognized its utility.

A similar point applies to sociobiological explanations of human behavior. Such explanations argue that common human behaviors stem from genetic predispositions. In effect, they argue that we have similar behaviors because we are descendants of the same group of human beings and hence share common genes through a line of genetic descent. But precisely because human beings are able to adapt themselves to the problems they face and pass these solutions on to others in the form of culture, we cannot necessarily infer that any particular set of behaviors stems from a line of common genetic descent. As we have seen, similarity of behavior across cultures may be due to common me-

metic descent, that is, cultural transmission. Or it may be due to the fact that two different cultures "invented the wheel" independently because they faced similar problems and devised similar solutions. In such cases there is neither common genetic nor common memetic descent.

Although many commonalities in human behavior surely do stem from our common genetic heritage, genetic descent is not the best explanation for large segments of common human behaviors. As Dennett points out, "In every culture known to anthropologists, the hunters throw their spears pointy-end first, but this obviously doesn't establish that there is a pointy-end first gene that approaches fixation in our species." People throw their spears in this way not because they are biologically programmed to do so, but because it makes sense to do so, and so everybody ends up doing it in pretty much the same way. A similar analysis applies to less frivolous examples, like the development of common systems of land tenure or accident law that appear in different times and places. The human condition often leads to similar problems across different environments; hence human reason produces similar behaviors to solve these problems; but it does not follow that the behaviors themselves are genetically predetermined.

Cultural Separation and Speciation

I noted earlier that cultural transmission is not simply a means by which memes are copied from one mind to another; it is also an important source of mutuation and change. Because perfect copying is the exception rather than the rule in memetic transmission, people's cultural software may vary considerably unless there are institutions and practices that homogenize it. Put another way, successful complexes of memes must have ways of accurately reproducing themselves in succeeding generations of minds if they are to survive. In fact, there are many devices for instilling common values and tools of understanding among members of a culture. The most simple is the existence of a common language, but we might also include the family, public schools, intellectual disciplines, and religious institutions.⁴⁸ These institutions have many different purposes. From an evolutionary perspective, however, they have one additional purpose: to preserve cultural content and cultural identity. They exist in order to reproduce memes (and hence themselves) in new minds.

Constant communication and participation in common social activities are important ways to reproduce and reinforce cultural software in the members of a culture. Conversely, isolation of individuals from a larger group results in cultural isolation and divergent cultural development. There is a useful analogy in evolutionary theory. Ernst Mayr argued that different species form because breeding populations become reproductively isolated, either because of geo-

graphic separation or because each inhabits a distinct ecological niche. This causes the genetic pool in the distinct populations gradually to diverge over time.⁴⁹

In like fashion, communicative isolation separates populations of memes, and over time these populations may develop in distinctly different ways. Communicative separation robs institutions of one of their most important means for memetic replication and cultural homogenization. Linguists have long understood that languages begin to differ from each other because of geographic isolation. Even cooking styles become distinctive when cultures are isolated.⁵⁰

Biological speciation results from separation of breeding populations, preventing genes from moving from one group to the other. Cultural speciation results from communicative separation, which prevents memes from traveling from the minds of one group to the minds of the other. This communicative separation may be geographical or spatial. But it may also be produced by culturally created boundaries that discourage communication between people and are themselves the product of previous cultural development. Thus if people who live next to each other never talk to each other—because cultural mores keep them apart—they may develop completely different ways of understanding the world. Under the right conditions, cultural differentiation can snowball—racial ideologies may keep blacks and whites from intermingling and communicating with each other, for example, leading to the development of increasingly distinctive subcultures and mutual incomphrension.

Disciplinary boundaries in the modern university exemplify another form of cultural separation. Disciplines are not only distinctive ways of thinking about things; they also serve as ecological niches that separate populations and produce divergent development. But instead of an ecology formed by the natural environment and other animals, this ecology is formed by other memes and cultural institutions. Other examples are clubs and societies that share common interests and develop their own distinctive preoccupations and languages.

Scholars who move across disciplinary boundaries often discover mutual incomprehension among members of different disciplines; each possesses a different vocabulary and different interests, research paradigms, and conceptions of what is interesting or important. As a result, an economist may find it much easier to understand a fellow economist three thousand miles away than the anthropologist in the building two blocks away.

Just as communicative isolation may tend to produce divergence in development, common experience and common communication may tend to homogenize the tools of cultural understanding in a population. Increasing communicative interaction can encourage reciprocal influence and shared ways of thinking. One must use the term *reciprocal* advisedly, though. The most numerous or dominant groups of individuals may have a disproportionate effect

on the cultural software of smaller and subordinate groups—unless, of course, the latter groups have greater communicative power.

This relation between commonality and cultural homogeneity suggests the signal importance of the rise of mass communication. Mass communication makes possible—indeed, increasingly enforces—enormous amounts of interaction between otherwise widely separated individuals and cultures. Much more than individual travel, mass communication is the great arbitrageur of cultural differences. It mixes cultural influences in ways that often annoy cultural purists. Moreover, because it multiplies opportunities for transmitting memes, mass communication also tends to accelerate the growth and mutation of forms of cultural understanding. Nevertheless, mass communication does not necessarily enforce uniformity; it simply creates more opportunities for mixing and reciprocal influence. Sometimes this mixing does produces homogeneity and uniformity, but sometimes it produces diversity and specialization.

Thus communication performs two contrary functions. On the one hand, it preserves stability and similarity between the various copies of cultural software located in each individual. On the other hand, it allows innovations in the tools of understanding to be transmitted to others, so that they may become part of the meme pool, the common cultural heritage. Communication is a source of stability as well as change in a meme pool and in the cultural software of individuals within a culture.

The Economy of Cultural Software

This book has offered two different accounts of the spread and development of cultural software. The first is conceptual bricolage: a non-Darwinian process of historical development through which human beings fashion new tools of understanding out of older ones, often with unexpected consequences. The second is memetic evolution: a Darwinian process of variation, reproduction, and differential survival of memes that form the building blocks of human cultural software. The first perspective describes the development and spread of culture from the standpoint of human thought, design, and action. The second describes this process from the standpoint of units of cultural transmission that compete for survival in the environment of human thought, design, and action.

We can view the spread and development of cultural software in a third way. We can see it as an economy of human communication—a process of exchange and development in which the members of a culture continually rewrite and reshape each other's cultural software. The idea of an economy joins the first two perspectives together, for it is both the mode of transmission of the products of cultural bricolage and the method of reproduction for the

memes that inhabit human minds. Equally important, the economy of cultural software is the means through which ideological power is wielded over members of a culture.

In accord with the computer metaphor, one might compare culture to a giant network of individuals. But culture is not a top-down network, in which a single server transmits identical copies of a software upgrade to the various nodes. It is more like the network of networks called the Internet, which has no center and in which an astonishing array of diverse information flows to and from different points simultaneously. Cultural software is not created in a single place, nor is it distributed from a central location, nor do all individuals share identical copies. The cultural software of individuals in a culture is written and rewritten through acts of communication and understanding among individuals in a culture. An individual's cultural software can also be rewritten through individual experience outside of interpersonal interaction. But the memes so created do not become cultural—in the sense of widely shared—unless they are transmitted to others. Hence even individual innovation and trial-by-error learning become part of the economy of cultural software through communication.

The nodes of a cultural network are continually communicating with and attempting to understand each other, and thus continually having reciprocal effects on the structure and content of each other's cultural software. This continuing process of communication is the economy of cultural software. Like other economies, it involves exchange, and it is driven by and operates through similarity and difference. Communication to others produces or reinforces homogeneity, even as differences in the understanding of individuals, however minute, are a potential source of change.

Although each individual has different cultural software, we can speak of "our" cultural software or the cultural software of a particular culture in two different ways. First, just as we can speak of a gene pool—the set of available genes that compete in the environment—we can also speak of a "meme pool." The meme pool of a given culture includes the copies of all memes that exist at any one time in the environment of human minds and information storage technology within the culture. Second, we can speak of this meme pool in dynamic terms—as an ongoing economy of transmission and exchange. This economy is the process through which the meme pool grows, develops, and is sustained. It creates the environment in which memes live and die, thrive and become extinct. The *economy* of cultural software is also the *ecology* for the memes that constitute individuals' cultural software.

When we speak of cultural software, we can either be speaking of the distinct collection of memes that forms part of a particular individual or of the larger economy of cultural software existing within a culture. But when we

speak of the cultural software of an entire culture, we must not think that we are describing a single great "program" that exists over and above each individual, or even a set of identical copies of a single program installed in isolated individuals. The cultural software of a group is not a separate set of skills in and of itself; it is rather a system of similarities and differences among the skills available to the members of a given culture. Both the similarities (which are sources of shared understandings) and the differences (which are sources of dissensus) are equally important parts of the economy. This economy is a huge system of networks, and networks of networks, of individuals continually communicating with each other by word and deed, by voice and action, continually engaged in a process of collective writing and rewriting of their cultural software.

Each person contributes to this economy through her words and actions, because she sends memes out into the world, where they can be absorbed and assimilated by others. Each individual is a potential source and a potential target of memetic infection. Through a partly cooperative and partly agonistic process, our tools of understanding are crafted and recrafted over time. This process produces a wide array of cultural skills, which are the collective property of the culture and are passed along to succeeding generations.

This set of available tools of understanding is the meme pool. It is sustained and replenished through acts of communication, just as the gene pool is sustained through reproduction. Through cultural transmission, each generation bequeaths to the next a huge collection of cultural skills, associations, heuristics, metaphors, conceptions, and constructs—a patrimony that will be squandered without perpetual communication between members of the culture.

Yet repeated transmission is also the source of change. Although symbolic and informational exchange is occurring all the time, there is no reason to think that it produces complete uniformity; indeed, it would be surprising if it did so. Communication continually introduces variation. Each person in the culture is equipped with slightly different tools of understanding and therefore carries away different experiences from communication. Each articulation of a meme in new contexts produces differences, however slight. Personal experiences and innovations of individuals give birth to new memes that join the meme pool once they are communicated. In this way, differences multiply over time, leading not only to the perpetuation of cultural software but also to its perpetual differentiation.

Consider, for example, the effects of rapid technological change on persons of different ages within a culture. Younger generations easily pick up technological skills and abilities that are difficult for older members to master, just as they develop linguistic habits and even accents that differ from their elders'. In the same way, we should expect that although the cultural software of each

individual overlaps with others in important ways, it also varies significantly as well. If enough people have cultural software that is sufficiently similar, this produces a cultural intersubjectivity that is also a cultural objectivity, because all of them see and understand the world in similar ways. This intersubjective agreement is accompanied, however, as it is in real life, by significant differences of understanding and belief.

Accounts of shared understandings usually face a problem in accounting for the dynamic nature of cultural traditions: How can a tradition grow and evolve while it remains a tradition shared by all of its members? How can shared meanings and practices remain shared if they are constantly changing? The twin concepts of the meme pool and the economy of cultural software allow us to give an account of this phenomenon. Shared understandings are the result of the partially similar (and partially different) cultural software of individuals within a particular culture. But this software does not remain the same indefinitely. Memes have differential rates of reproduction and survival in the environment of human minds and their technologies of information storage. This causes the cultural software in the minds of individuals to evolve. But as long as the members of the culture are part of the same meme pool and participate in the same economy of communication, their understandings evolve together in roughly the same way. Biological species continue to share a common gene pool and evolve together even though that gene pool is constantly evolving as members continue to interbreed. In the same way the economy of communication among members of a cultural tradition ensures that shared understandings continue to be shared by individuals even though the content of these understandings changes over time as the meme pool constantly changes.

In this way, the theory of cultural software offers a distinct improvement on historicist accounts of cultural understanding like Gadamer's. It translates the idea of a historically evolving tradition into something that truly exists in each individual and constitutes each individual. It shows that the tradition evolves as an economy of communication that regulates a shared meme pool. The theory thus avoids the theoretical puzzles that stem from supraindividual entities—like a tradition, a collective consciousness, or a Zeitgeist—offered to account for the commonality of beliefs and actions. The claim that there is a "spirit of the age" that produces similarities in artistic and intellectual production, for example, merely begs the question of what such an entity is, where it is located, and how it can have causal effects on individual thought and action.

In contrast, the theory of cultural software explains commonalities in intellectual and artistic production as the result of the similarities in the cultural software found in different individuals within a culture. These similarities are maintained by an economy of exchange, reproduction, and evolution. Thus, what people call collective consciousness or the spirit of the age is not a cause of similarities in individuals' cultural production; it is the apparent effect produced by an economy of exchange among people with sufficiently similar cultural software. Moreover, unlike these hypothetical entities, this system of exchange not only produces and reproduces relevant similarities among individuals; it also produces and reproduces differences that lead to divergence and variation. Thus we can say, without the introduction of any mysterious entities, that painters in the Renaissance or composers in the Classical period had similar styles not only because they used the same technologies of painting or music, but because they employed similar tools of understanding. In a given culture at a given time, individuals in different walks of life and different intellectual pursuits produce artifacts and theories that bear uncanny metaphorical similarities to each other because the tools that lie to hand in that age are similar for each of them, because each thinker draws from the same meme pool. We need not say that these similarities exist because of the Zeitgeist; rather we should say that the metaphor of the Zeitgeist describes the operation of an economy that produces these similarities.

The Distribution of Cultural Software

An economy of cultural software is a system of similarity and difference in the memes that constitute the members of a culture; the degrees of that similarity and difference may vary in different cultures. Hence an economy of cultural software is distinguished not only by the content but also by the distribution of different types of cultural software among its members. The relative distribution of similarity and difference affects the degree of intersubjective agreement in a culture, as well as the degree of disagreement, mistake, and dissensus.

The distribution of memes in a culture is an important feature of the ecology in which memes spread and evolve. If the distribution of memes changes in a culture, the character of the culture may change dramatically. Durkheim's notion of collective consciousness, for example, described the thought of relatively primitive societies. But this consciousness dissipated as these societies developed increasing specialization of labor and moved away from mechanical solidarity toward the organic solidarity that we associate with modernity.⁵¹ The dissolution of collective consciousness corresponds to a change in the distribution of memes as well as their content.

People often identify modernity with increasing secularization, rationalization, and differentiation of social functions. But we can also think about modernity in distributional terms. What distinguishes modern (and postmodern) cultures is more than the common possession of a particular set of tools of understanding—they also possess a more exaggerated and distinctive economy of differences in cultural software that, in turn, produces the kinds of relativism and historicism, disenchantment and lack of solidarity that we associate with modernity.

There is a familiar view of modern thought as the result of diverse cultural influences meeting in a single culture. This mixing of influences may stem from changes in communications technology, increased opportunities for travel or trade with other cultures, or increasing rates of literacy and education. In memetic terms, all of these tend to flood the existing meme pool with memes from other populations. This memetic invasion tends to change the distribution of the pool. The predictable result is wider disparities in cultural software as well as mixing and crossing of cultural lineages.

Changes in distribution also effect changes in content. First, old memes tend to mix with new ones, spurring cultural innovation. Second, particular memes and memetic filters proliferate in response to the flood of new memes. Some of these are the familiar tropes of modernist anxiety—a sense of loss of an organic connection to past traditions, a desire to regain cultural authenticity, the longing for an imagined golden age of uncomplicated consensus and harmony, and the fervent need to regain the past by clinging to its symbols and material manifestations.⁵² Another very different set of memes also flourishes in this new ecology—memes that promote cultural relativism and skepticism. The ecology of modernity is a fertile breeding ground for these ideas because the very presence of so many different and conflicting cultural influences seems to provide evidence for them.

The past two chapters have portrayed cultural understanding as a result of an ongoing economy of communication through which individuals transmit memes to one another and rewrite one another's cultural software. Implicit in this picture are deep connections between cultural communication and ideological power. Communication is a potential source of power over other individuals because it can rewrite their cultural software. Conversely, our ability to understand others is a potential source of vulnerability, because it means that we are susceptible to ever new forms of memetic invasion.

This connection between power and cultural understanding brings us back to the theory of ideology. In the next three chapters, I shall explain how the theory of cultural software approaches the traditional questions that have been asked about ideology and grapples with the recurrent problems that any theory of ideology must face.

There are many different definitions of the concept *ideology*, and many different ways of approaching its study. In particular, a theory of ideology must consider the following questions:

- 1. What kinds of things (objects, entities, mechanisms, or structures) are we investigating? This is the problem of the proper object of study.
- 2. Do we define ideology in terms of its content (for example, distortion or mystification), the functions it serves (for example, furthering the interests of the ruling class), its causes (for example, cognitive bias, reduction of cognitive dissonance), or its effects (for example, creating or sustaining unjust relations of social power)? This is the problem of the proper mode of explanation.
- 3. What is our attitude toward ideology—pejorative, positive, or neutral? This is the problem of interpretative stance.
- 4. How does our theory handle the inevitable difficulty that the analysis of ideology may itself be ideological? This is the problem of self-reference.

The first two questions are the subject of the present chapter; the last two are the subject of Chapter 6.

The Object of Study

Some theories of ideology define their subject matter in terms of beliefs, held either by groups or by individuals. Jon Elster, for example, defines ideology as false or distorted conscious beliefs held by individuals about the social world. Other theories of ideology are concerned with linguistic or cultural products

or social practices of meaning that do not necessarily exist in the minds of individuals but that individuals use in understanding the social world. Thus John Thompson and Clifford Geertz view the study of ideology as the study of "symbolic forms." This expression includes such diverse phenomena as linguistic utterances, symbols, literature, traffic signs, television broadcasts, and advertisements. In Geertz's case this choice is a deliberate attempt to externalize the object of study—to move from the investigation of internal mental processes to external observable entities like written symbols, linguistic utterances, artistic objects, and behavioral practices.

The theory of cultural software takes as its object of study tools of human understanding produced by cultural evolution. Symbolic forms play a key role in cultural evolution because they carry units of cultural transmission; hence the study of symbolic forms is crucial to the study of cultural software. When we study a symbolic form (such as an advertisement), however, we are interested in the ways of understanding that produced it and the effects that it has on the way that others understand the world. Hence we are interested in cultural artifacts and symbolic forms for four reasons. First, symbolic forms are effects of cultural software and therefore are evidence of the mechanisms of thought. Second, symbolic forms have reciprocal effects on individual cultural software. Third, symbolic forms are media through which minds communicate and by which they share meaning. Fourth, symbolic forms are a common terrain of negotiation and struggle over shared meanings. These negotiations and struggles, in turn, affect the cultural software of the individuals who engage in them.

Moreover, the object of our study is necessarily broader than conscious beliefs, at least if by this term we mean beliefs that can be expressed in the form of propositions, like "Jews are greedy," "Women don't make good pilots," or "The Holocaust never happened." To be sure, beliefs can be tools of understanding and can be used to create new tools. But more important objects of study are cognitive mechanisms that produce beliefs. Examples include the tendency to structure experience in terms of narratives, psychological methods of categorization, varieties of metaphoric and metonymic thinking, strategies for reduction of cognitive dissonance, heuristics and biases employed in making judgments under uncertainty, and understanding by means of networks of conceptual oppositions in the form "A is to B as C is to D." Propositional beliefs can be true or false, but cognitive mechanisms are neither true nor false. Rather, they are the ways in which attitudes and judgments are formed: they produce beliefs that can be true or false.

For example, consider the tendency, noted by many feminist writers, for people to think of the male as the standard case or unspoken norm of gender, so that the feminine is treated as an afterthought, an additional feature, or a special case. 4 This tendency is produced by various mechanisms of understanding, some of which will be discussed in more detail in Chapters 10 and 11. Nevertheless, it is important not to confuse these mechanisms with the propositional belief "Men are the normal case and women are the exceptional case," or the directive "Think first in terms of men and then consider women later as an afterthought if it is brought to your attention." Individuals' practices of thought may be aptly summarized by such a propositional belief or such a directive, but this does not mean that the mind consciously employs such beliefs or rules in forming its judgments. Social understanding does not proceed exclusively or even predominantly at the level of such conscious propositional belief or conscious rule following. Understanding also occurs through various mechanisms of framing, narrative construction, characterization, and categorization. Nevertheless, like the study of symbolic forms, the study of conscious beliefs is important because through them we can attempt to understand the mechanisms of social understanding that produce them.

The Focus of Study

Once a theory of ideology has settled upon its objects of investigation, it can study these objects in many different ways. We can study them in terms of their content, their causes, their effects, or the social functions they serve. Often the way that ideology is defined leads to a focus on some of these aspects to the detriment of others. If we define ideology purely in terms of false or distorted beliefs, for example, the study of ideology becomes centered on the question of the content of beliefs. To study ideology is to study how certain beliefs are false or misleading; hence when the analyst has revealed this falsity or distortion, her task is largely completed.

Marxist theories of ideology often approach ideology in functional terms. They study how ideologies serve the interests of various classes. Michéle Barrett summarizes the classical Marxist definition of ideology as "mystification that serves class interest."5 This definition is functional (although it might be restated in nonfunctional terms): ideologies are defined and studied in terms of the social interests they serve. A functional approach, however, is necessarily limited. It tends to explain the development and content of ideologies solely in terms of the class interests that they further, rather than offering evolutionary or other causal mechanisms that explain how and why ideological beliefs are produced. Put another way, functional accounts cannot reliably serve as causal explanations: even if a belief serves the interests of a particular class, it does not follow that the belief was the result of something that class did. Similarly, showing that something serves the interests of one class does not by itself explain how a belief was generated or held by another class.6

In contrast, Jon Elster rejects functional approaches to ideology entirely.⁷ His theory seeks to offer purely causal explanations of ideology. Thus, although he defines ideologies in terms of their content as false or distorted beliefs, he is specifically interested in how these beliefs are caused by various social psychological mechanisms.⁸

Finally, we might define ideology in terms of its effects. For example, John Thompson defines the study of ideology as the study of how symbolic forms create or sustain relations of domination. Thus his approach is centered on what ideology does rather than what causes it or what interests it serves. In Thompson's view, approaches that focus on content are insufficient because the content of a particular symbolic form does not by itself tell us whether it helps sustain relations of domination.⁹

The theory of cultural software has a twofold focus: First, it is concerned with how tools of understanding are produced through conceptual bricolage. Second, it is concerned with how these tools of understanding help create or sustain injustices in particular social contexts. The goal of this theory is not primarily functional explanation but causal or evolutionary explanation. It studies the tools of understanding in terms of the causes that produce them and the effects that they in turn produce. Although this study is obviously concerned with the content of beliefs and symbolic forms, that inquiry is subsidiary to the study of the effects produced. Moreover, because the theory focuses on just and unjust effects, its analysis is overtly normative as opposed to merely descriptive.

Under this approach there is, strictly speaking, no longer a single thing called ideology. The theory of cultural software, while dissolving the study of ideology into the larger study of cultural understanding, also breaks the study of ideology down into the study of ideological mechanisms and ideological effects. Ideological mechanisms are mechanisms of social cognition that produce ideological effects. Ideological effects are effects of cultural software that help create or sustain unjust social conditions, unjust social relations, or the unjust use of social power. Ideological thinking, in short, is employment of ideological mechanisms of cultural software that produce ideological effects. Note that symbolic forms produced through the use of cultural software can also have ideological effects through their effects on human understanding. For example, perfume advertisements can have ideological effects if they help to create or sustain unjust relations between men and women.

The phrase "help to create or sustain" in the definition of ideological effects must be understood in a limited way, for otherwise the definition is seriously overinclusive. If a person used statistical formulas to calculate the numbers of individuals who could be transferred to a concentration camp, we would not say that the mere skill involved in applying the algorithm was an example of

ideological thinking, even though it would literally be a use of cultural software (a mathematical skill) that helped maintain unjust social conditions. Rather, cultural software has ideological effects when it creates ways of thinking about the social world or about others in the social world. Although the study of ideology is not concerned with the skills involved in statistical computation, it is concerned with the ways of looking at people that lead to the judgments that it is appropriate to apply these statistical methods to facilitate genocide. These judgments include, among others, that people should be shipped to concentration camps because they are inferior or that it is appropriate to think about people as commodities that must be efficiently shipped to the most efficient locations for the most efficient forms of slaughter.

These definitions of ideological effects and ideological mechanisms make what is ideological turn heavily on social context. Cultural software has ideological effects only when and only to the extent that it results in various forms of injustice. This means that in other contexts cultural software may have no significant ideological effects. Moreover, even when cultural software has ideological effects, these effects do not exhaust its social meaning, its content, or its usefulness. We can make a similar point about symbolic forms: a perfume advertisement is not merely a symbolic form with ideological effects; it is also, among other things, an advertisement for perfume. More generally, the tools of cultural understanding may have many other features and advantages and may serve many other functions apart from their tendency to produce ideological effects in certain circumstances.

In like fashion, ideological mechanisms are defined contextually. When mechanisms of social cognition produce ideological effects, one can speak of them—for this purpose and to this extent—as ideological mechanisms. But they are ideological not because of their inherent nature but because of the context in which they are employed and the effects that they have.

What distinguishes ideological thinking from mere fantasy or mistake is the social context in which belief occurs and the use that people make of it. An important consequence of this approach is its emphasis on the normative dimension of all ideological analysis. To understand what is ideological, we need a notion not only of what is true but also of what is just. False beliefs about other people, no matter how mistaken or unflattering, are not ideological until we can demonstrate that they have ideological effects in the social world. To demonstrate this, we must know something about the relationship between a person's thought and the existing conditions of social power, as these provide the necessary background for considering questions of justice and injustice.

For this reason, the study of ideology necessarily intersects with the study of how social power is created, sustained, and distributed, because one of the objects of this study, the ideological effect, is a highly contextual product of cognitive capability and social situation. To be sure, sometimes we may infer that a particular way of thinking—a white American's belief that all black people are lazy and immoral, for example—is so likely to produce or contribute to injustice that we may consider it presumptively ideological. But this is the case only because we already understand the social context in which this way of thinking occurs, the forms of behavior it is likely to lead to, and its place in a larger social system of race relations. Nevertheless, unjust social relations or unjust social power may be created or sustained in many different ways that are not always easily discernible from the content of a particular belief, especially when the context is unusual or unfamiliar.

Perhaps the best example of this principle is the bizarre phenomenon of Japanese anti-Semitism. There are very few Jews in Japan today and thus very few opportunities for discrimination against them. Nevertheless, anti-Semitic books and comments have appeared continually in Japan over the years, often repeating the most vicious claims of Nazi ideology and Eastern European anti-Semitism. ¹⁰ Especially popular are beliefs about a secret worldwide Jewish financial and media conspiracy of enormous scope and power. What is most amazing is that the very same libels that in the European context were part and parcel of a terrible social system of discrimination (and extermination) are combined in Japan with a peculiar form of philo-Semitism in which Jews are admired for their supposed shrewdness and business acumen. ¹¹

All of this is not to claim that Japanese anti-Semitism has no ideological effects. Rather, my point is that we must not conflate this phenomenon with European anti-Semitism even though its beliefs and slogans appear to be similar in content and may even have their origins in European anti-Semitic literature. The ideological effects of Japanese attitudes toward Jews seem to have more to do with supporting and sustaining a larger system of beliefs about business and economic competition in Japan. These ways of thinking, in turn, may help sustain relations of unjust power not between the Japanese and a Jewish minority but within Japanese society itself, or between the Japanese and the outside world. Moreover, Japanese anti-Semitism also serves as a way of expressing anti-American sentiments, which have surfaced as Japan and the United States have increasingly become economic adversaries. Because Jews are portrayed as the hidden masters of American business and government, anti-Semitic rhetoric becomes another way of complaining about American culture and American trade policies.¹²

Of course, if large numbers of Jews were to emigrate to Japan, existing anti-Semitic attitudes might lead to unjust treatment of Jews, just as they did in Europe and America. This is yet another consequence of my basic point about the uses and effects of conceptual tools. When introduced into new social settings, the tools of understanding display different effects, benefits, and disadvantages. That is why the study of ideology cannot rest on content alone but must take into account the environment in which cultural software operates. Indeed, the view that the power of ideas lies in their content and not in their content in a particular context is itself a way of thinking that causes us to misunderstand social situations.

This approach sheds a somewhat different light on so-called beneficial ideologies. Suppose that the students in a particular elementary school classroom are falsely told that they are very bright and very able, indeed, much more bright and able than other students of their age. As a result, their test scores, as a group, actually begin to improve. The source of their esprit de corps is fraudulent, yet it seems to benefit them. This is an example of what Jon Elster calls the "benefits of bias." Such situations are important to explain in the Marxist tradition because ideology is often defined functionally in terms of what serves the interests of a particular class. Hence it follows that some ideologies, while false, may actually benefit the people who hold them—for example, the members of the bourgeoisie. Because I define ideology in terms of what is just rather than what is in a particular group's interests, my analysis of this example is quite different: we cannot yet even say that these students are engaged in ideological thinking until we study how justice might be affected by their views of themselves. First, these students may start to look down on students in other classes and other schools and to discriminate against them although the others have equal or greater abilities. Second, some students in the class may not be able to live up to their teacher's claims of superior ability, and they may engage in strategies of dissonance reduction to avoid this recognition: for example, they may be more likely to assume that people who criticize their work are simply mistaken, or they may come to think that their failures are due less to ability than to luck or sheer coincidence. This may harm them in the long run. Thus, a so-called ideology of superior achievement is not ideological thinking in my sense of the word unless and until it has particular effects, and then only to this extent. Moreover, the flip side of this claim is that even the most seemingly benign and beneficial forms of thinking can have unexpected and unfortunate effects as they are extended into new contexts and situations. It is at that point that they become forms of ideological thought.

Above all, this approach does not view ideology as something separate from cultural understanding. The mechanisms of what we call ideological thinking are no different in kind from the ordinary forms of thought. There is not a separate set of devices that constitute "the ideological" and another set that constitute "the nonideological." There are not mechanisms of social cognition that always produce ideological effects and other mechanisms that never do so. In particular, we must resist the natural tendency to think that ideology constitutes a separate, deviant form of social cognition that can readily be distinguished in terms of its operations from the supposedly normal, nonideological forms and mechanisms of thought that characterize everyday reasoning. The mechanisms of ideology are the mechanisms of everyday thought, which in particular contexts produce effects that are both unfortunate and unjust. Conversely, the mechanisms of everyday thought can become ideological mechanisms if they are employed in inappropriate contexts and situations.

This conclusion is consistent with our earlier discussion of conceptual bricolage. The tools of our understanding can be alternatively advantageous and disadvantageous as they are applied in new situations and new contexts. Among the many possible disadvantages that conceptual tools can have is their tendency to promote injustices; conversely, one of their many possible advantages is the relative lack of this tendency. Thus tools of understanding that are entirely benign in some circumstances may become malignant if too much is demanded of them or if the context in which they are employed changes sufficiently. Then their limitations become apparent in the same way that many other disadvantages of tools may suddenly surface.

The temptation to identify ideology with a sort of pathology may stem from the familiar notion that ideology is false or distorted belief. Given this assumption, it seems natural to think of falsity or distortion as a kind of illness or malady, especially if it has harmful effects. For example, we often speak of racism or anti-Semitism as a sickness or a disease. In fact, the metaphor of disease is not completely unreasonable, as I shall discuss momentarily. But identifying ideology with pathology simply because beliefs are false or distorted improperly focuses on content rather than mechanisms—or, to use the metaphor of disease, it focuses on symptoms rather than etiology or cause. From the standpoint of causal mechanisms, the question is whether the effects that people have traditionally assigned to the ideological are due to (1) a special mechanism different from the ordinary mechanisms of social cognition; (2) the extension or employment of cognitive mechanisms into contexts for which they are not well adapted; (3) a spontaneous malfunction in cognitive processes; or (4) the invasion of some external force into normally and properly functioning cognitive processes that causes them to malfunction. I reject (1) and suggest that many ideological effects are produced by (2).

This leaves cases (3) and (4), both of which explain ideological effects in terms of malfunctions. Obviously, there is some overlap between the notion of overextension and the notion of malfunction. Nevertheless, the concepts are not identical: we would not say of an airplane that it malfunctions because it is a poor vehicle for traveling on land. This is not malfunction but maladaptation. One could collapse the distinction between malfunction and maladaptation only if one assumed that our tools of understanding should be capable of understanding everything in all contexts. Then to the extent that they failed

to do so, we would say that they were malfunctioning. This seems to ask too much of our tools of understanding, though; after all, no tool exists that is equally well adapted to all tasks.

Much of the distortion that we see in ideology involves the side effects of tools of understanding that become prominent and maladaptive in particular contexts. Ideological effects are usually the unexpected and unpleasant side effects of conceptual bricolage. I do not reject out of hand the possibility that some ideological effects are due to a genuine malfunction in cognitive processes. But this malfunction would have to appear in many individuals at once in order to qualify as an ideological phenomenon. A simultaneous malfunction by members of a culture is unlikely. This leaves the possibility that if some ideological phenomena are due to a malfunction in our cognitive processes, it is a malfunction brought on by some external force that affects many people at once. One possibility is that when individuals are placed in situations with which their cognitive systems cannot cope, they break down or malfunction, just as we say that a car malfunctions when it is forced to drive through water, or a vacuum cleaner malfunctions when it is forced to deal with too great a quantity of dust. If many individuals face the same type of experience, this malfunction would be similar for all of them. But it is hard to imagine that this explains most ideological effects. After all, human intelligence is quite adaptable, and many ideological effects, like racism or anti-Semitism, are longlasting phenomena that occur over many generations. The idea of a long-term breakdown in cognitive processes seems implausible.

Instead, the theory of cultural software offers a somewhat different account of how relatively robust and long-term ideological effects can be produced by a malfunction due to an "external entity." This external entity is none other than cultural software itself, transmitted from other individuals and spread throughout a culture like a computer virus. A computer virus is just a special kind of computer software that is able to spread and reproduce itself in other computers. By analogy, cultural software may act like an informational virus that infects one node on a network and then, through the exchange of information, gradually infects all the others.

Under this model, long-lasting and widespread ideological effects are produced by informational or cognitive "viruses" that are passed from person to person and generation to generation. If so, we might think of racism or anti-Semitism as a sort of socially spread informational virus or parasite that, while not totally debilitating subjects, affects their behavior and cognition for the worse.

In fact, this model of ideological effects is the model of memetic evolution through cultural communication. Memes are reproduced in individuals through a social network of communication and transmission. The spread of ideological

viruses is merely a special case of the basic mechanism through which cultural software is written, transmitted, and modified. All cultural software can be thought of as a kind of informational virus, transmitted from person to person; or, put another way, what we might call an ideological virus is just another kind of cultural software. Our devices for understanding the social world are constituted in large part by idea-programs that were able successfully to be transmitted to us and absorbed into our cultural software. The complexes of memes that give rise to racism and anti-Semitism, in this sense, are no different from any other idea-programs—like those producing predilections for free speech or free markets—that make use of our cognitive capacities to grow, spread, and develop, just as genes "use" bodies in an evolutionary system.

Hence what differentiates cultural software from a so-called ideological virus is the harmful effect that the latter produces in a particular social context. As the example of European versus Japanese anti-Semitism demonstrates, an ideological virus can produce very different effects when it is introduced into different environments. If an informational virus produces no such harmful effects—just as there are many viruses in the human body that are relatively benign or harmless—then it does not produce an ideological effect. Fantasies about people in far-off lands may be distorted and false, but they do not become ideological until there are conditions of justice between the two peoples—that is, until there is communication, trade, and the possibility of war, conflict, struggle, economic exploitation, or colonization. Then these fairy tales (which may already have had certain ideological effects within a culture) take on a more serious and harmful tone. Fantasy becomes ideology when justice is at stake.

This line of reasoning brings us back to our original hypothesis—that ideological effects are produced by ordinary mechanisms of thought that have harmful or maladaptive consequences in particular contexts and situations. Ideological effects occur when cultural software "goes wrong" in some important way. The power of ideology over our imaginations is a special case of the power that all cultural software has over our imaginations.

The power of ideology within this picture is quite different from the picture underlying a more traditional Marxist theory of ideology. In the traditional account, ideas have power because they present a distorted picture of reality to the minds of the persons holding them, causing these persons to act against their objective interests. From the standpoint of the theory of cultural software, the power of ideology is the power of the culturally produced capacities of our minds to shape social reality for us, and thus simultaneously to empower and to limit our imaginations.

This approach makes considerable use of concepts like usefulness, ade-

quacy, and suitability. But these concepts can hardly be considered inherent properties of the tools of understanding. Adaptability is a judgment made about the operation of a tool in a particular context. It is also a judgment made by an observer who assesses the operations and effects of mechanisms of thought. This means, among other things, that the study of ideology is necessarily an interpretive endeavor, although this fact makes it no less useful. Finally, because all cultural and social understanding makes use of cultural software, all ideological analysis—that is, all judgments about the existence and nature of ideological effects—involves judgments by an analyst that employ the analyst's cultural software. This raises problems of self-reference, which are discussed more fully in the next chapter.

What Kinds of Effects: Hegemony or Unjust Power?

The study of ideology necessarily has a normative dimension. It cannot be value free but must presuppose a view about what is good and bad, advantageous and disadvantageous, just and unjust. The analyst cannot describe and analyze ideological effects without reference to concepts like truth or justice. She must make interpretive judgments about what social conditions are like, and she must also make judgments about whether a way of thinking is adequate or inadequate to serve particular ends and whether social conditions are just or unjust. Ideological analysis does not end with a demonstration that a particular belief or symbolic form is partly or wholly false or distorted. It must ask how this falsity or distortion might create or sustain unjust social conditions or unjust relations of social power. Thus ideological analysis does not merely involve considerations of truth and justice; it is fundamentally a question of the relationship of truth to justice.

Because I define ideological effects in terms of actual or potential injustices rather than the presence of hegemony or domination, it may be helpful to contrast my approach with that recently offered by John Thompson. Thompson defines the study of ideology as the study of how symbolic forms create or sustain conditions of domination. He then defines domination in terms of systematic asymmetries in relations of power—that is, "when particular agents or groups of agents are endowed with power in a durable way which excludes, and to some significant degree remains inaccessible to, other agents or groups of agents, irrespective of the basis on which such exclusion is carried out." Under Thompson's definition, women in the United States would be dominated if we could show that they are disadvantaged vis à vis men systematically in many different ways, including jobs, income, status, education, economic opportunities, and other resources. Thus while Thompson argues that the es-

sential feature of ideology is the creation or preservation of domination, I have argued that it is the creation or preservation of unjust power or unjust social conditions.

One reason for this difference is that Thompson's definition is underinclusive. Not every example of ideological thinking contributes to systematic asymmetries in social resources or power relations between groups. Consider, for example, the phenomenon of black anti-Semitism in the United States. Anti-Semitic propaganda by black nationalist groups like the Nation of Islam does not contribute to or produce systematic asymmetries in resources or power relations between blacks and Jews or even between all Christians and Jews. Indeed, in contrast to blacks, Jews have been relatively successful in gaining access to social resources in the United States. For this and other reasons, Jews provide a convenient scapegoat for some members of the black underclass, just as blacks themselves have provided a convenient scapegoat for lower-class whites in the United States. Black anti-Semitism, like resentment and hostility among some blacks toward Asian Americans, is in part the result of competition between various minority groups; it is not a means by which blacks oppress Jews or Asians and systematically deny them access to social resources. Nevertheless, anti-Semitism and anti-Asian beliefs may in fact lead to particular injustices—acts of violence, for example—against Jews or Asians, either by blacks or by other groups. Thus a focus on systematic asymmetries in power defines ideology too narrowly; the study of ideology must be concerned with injustices produced by tools of understanding whether or not they stem from domination of a subordinated group by a dominant group.

To be sure, black anti-Semitism or anti-Asian sentiments may also contribute to the perpetuation of systematic asymmetries between blacks and whites, by diverting attention onto scapegoats and away from positive solutions to the challenges the black community faces. Similarly, prejudice against other racial minorities, like Asians or Hispanics, alienates potential allies who might otherwise fight together with blacks against white supremacy. Nevertheless, the ideological effects of black anti-Semitism or anti-Asian prejudice are not exhausted by their ability to hinder black economic progress and further white supremacy. Even if these prejudices did not harm the just interests of blacks, they would still be ideological, because they can and do lead to injustices between members of different minority groups.

Thompson's formulation suffers from these difficulties because it has not yet thrown off the shackles of a traditional Marxist model that envisions a dominant class, a subordinate class, and an ideology that justifies the subordination of the latter by the former. Systematic group domination by a dominated class over a subordinate class is the central concern; it follows that forms of social injustice or unjust social power that do not involve hegemony are not

properly the concern of the theory of ideology. As Thompson himself stresses, "Ideology, according to this conception, is by nature hegemonic, in the sense that it naturally serves to establish and sustain relations of domination and thereby to reproduce a social order which favors dominant individuals and groups." Hence Thompson limits ideology to the study of "the ways in which meaning is mobilized in the service of dominant individuals and groups."15

Unfortunately, this model is too simplistic to describe a large number of ideological phenomena, particularly in a country like the United States, where there are many different groups with varying degrees of social power and multiple and cross-cutting social identities. Antiblack prejudice by Korean Americans and anti-Korean prejudice by American blacks cannot easily be subsumed within a hegemonic conception of ideology. Nor does a hegemonic approach contemplate the possibility of simultaneous membership in groups that are dominant and subordinate—working-class white males who are homosexual, for example, or upper-class heterosexual women who are not physically disabled. One is perfectly free to limit the scope of the study of ideology in this way, of course, but the danger is that a large portion of what most people would consider ideological phenomena will be missed. Moreover, this limitation may have significant ideological effects on the analyst's own thought about ideology and social conditions.

For the theory of cultural software, the equation of ideology and hegemony is problematic for seven additional reasons. The first stems from the basic point that ideological mechanisms are the mechanisms of everyday thought about the social world. There is no reason to think that the kinds of cognitive mechanisms producing ideological effects that benefit dominant groups and harm subordinate groups are different in kind from those producing benefit and harm to other groups. It is likely that the mechanisms that produce prejudices between groups are fairly similar, although the results may differ because of the relative positions and histories of various groups in society. If we restrict our study of ideology to mechanisms producing beliefs that benefit dominant groups, we cut ourselves off from many examples of ideological thinking that not only shed considerable light on more hegemonic examples but are fully worth studying in their own right.16

Second, the concept of a dominant ideology leads us to view ideology in terms that are too monolithic. What people usually think of as ideology is really the confluence of many different types of cognitive mechanisms. The ideology of patriarchy, for example, is not a single thing, or a coherent system, but rather a group of heterogenous and partly reinforcing ideological effects. This heterogeneity may be one cause of its adaptability as well as a source of its possible deconstruction and subversion.

Indeed, there is a notable tendency among theorists of ideology to confuse

the pervasiveness and the wrongfulness of a worldview with its systematicity. Thus, Catharine MacKinnon, in a famous passage, has described patriarchy as a "metaphysically nearly perfect" system.¹⁷ This way of thinking may itself betray a certain ideological effect, because it conflates the powerful with the well-ordered. Ideologies in the larger sense that MacKinnon is concerned with are always the product of bricolage and memetic evolution. Hence they lack the characteristics of design: they always have conflicting and variegated elements, their seams always show, and loose threads are always dangling. Of course, this makes them no less powerful: an avalanche of motley elements is still an avalanche. But it does suggest that the theorist of ideology may be misled if she attempts to fit the entire phenomenon into a single, systematic analysis rather than looking for the confluence of various ideological effects and for their possible points of interaction and conflict. Indeed, the heterogeneity of cultural software is important precisely because it makes possible forms of resistance to received ways of thinking.

Third, the notion that ideology is concerned only with the preservation and maintenance of dominant ideologies neglects the importance of competition between various ways of thinking within a culture. This competition occurs at many different levels and at many different places in society; there are not simply two armies contending on the field, and those armies that do contend already are fragmented and partly divided against themselves. Within American society for example, many different and partially overlapping groups promote their ways of thinking about the social world; and many different currents and eddies of social power result from these encounters. Together these encounters produce heterogenous matrices of social power, mixing together the just and the unjust in an atrocious and unpalatable stew. To see only some elements of this mixture as worthy of the title of ideology is itself ideological, for it hinders the identification and critique of the many forms of social injustice that do not correspond to the grand narrative of the "hegemonic."

Fourth, when we define ideology in terms of symbolic forms that benefit dominant groups, we risk sentimentalizing the attitudes and interests of other groups, in particular subordinated groups. We risk overlooking the possibility that the beliefs of subordinated groups can also be distorted, self-serving, and unjust to other groups, even including more dominant groups. There is no reason to think that self-serving or distorted views of the social world are confined to dominant groups. Prejudice tends to beget prejudice and hate tends to beget hate. Persecution can lead to persecution complexes. Moreover, if a group's opportunities and access to knowledge have been limited by its social condition and its comparative lack of social power, this may seriously affect its members' understanding of the social world, producing ideological effects in their thought.

Even when subordinated groups have a relatively adequate understanding of the social world, it by no means follows that what these groups believe to be in their interest is always just, or even that what is actually in their interests is always just. This is especially so, one might think, in a multicultural society in which many different subordinated groups scramble for social betterment and political power. An obvious example involves tensions between black and Hispanic communities in the United States over the drawing of district boundaries that effectively determine the result of elections to state and federal legislatures. Black and Hispanic communities may correctly recognize that drawing boundary lines one way rather than another would guarantee the election of a black or a Hispanic representative in Congress, and they may also correctly assess that this would further black or Hispanic interests. Nevertheless, it is entirely possible that furthering one group's interests in this way may be unjust to other groups. We are no longer in the Marxist world, where furthering the interests of the oppressed (the proletariat) necessarily furthers justice or the proper direction of history. The belief that something is just simply because it favors a subordinated group may itself, under some conditions, involve ideological effects.¹⁸

Fifth, defining ideology in terms of what benefits members of dominant groups is problematic because subordination is not simply an on-off property of individuals or groups. There are different degrees and kinds of subordination among different groups, and individuals have multiple group identifications. Thus it is possible for an individual to be in a subordinated position with respect to some groups but in a privileged or dominant position with respect to others—consider the example of white middle-class heterosexual women.¹⁹ Pursuing the interests of white women may infringe on the just interests of black men, and vice versa. The endless possibilities for self-serving views of the social world between groups all of which can claim to be subordinated in one way or another-along with the concomitant injustices that may be produced by these views-shows how limited and simplistic a bipolar dominatordominated model can be, and demonstrates the need to expand the notion of ideology beyond a hegemonic conception.

The traditional proletariat-bourgeoisie model avoids these problems, first because it tends to reduce the number of groups to two, and second because it assumes that what is in the proletariat's interest is necessarily just or at least follows the course of proper historical development. Within this model the problem, rather, is ensuring that the proletariat understands what is in its own interest—that is, ensuring that it develops an appropriately revolutionary consciousness. Nevertheless, in a society where injustices do not derive wholly from economic power, in a society that features many competing and partially overlapping groups, divided on the basis of race, ethnicity, religion, language, gender, sexual orientation, and disability, one can no longer employ such simplifying assumptions. The simple model of dominator and dominated itself threatens to become ideological because it obscures the complexity of social conditions.

Not only does a bipolar approach tend to neglect the many different kinds of subordinated groups, it also tends to collapse, homogenize, and demonize the interests, attitudes, and beliefs of whatever group is described as dominant. Such a homogenization may disguise fragmentation within the dominant group as well as the existence of relatively subordinate and distinct subgroups. To speak about hegemony by whites, for example, is to forget that some women are also white; to speak about the hegemony of white males is to neglect the fact that the interests of lower-class white men may be quite different from those of more affluent white males. Moreover, the homogenization of white males into a single group obscures the fact that some of the most vitriolic race hatred appears not among the most powerful members of white society but among the most disaffected and disenfranchised. The Ku Klux Klan and other white supremacist organizations have often found that the poor and uneducated are more promising recruits than the well educated and the well-to-do. One reason why such groups turn to rabid racism, anti-Semitism, and anti-Catholicism is that, given their economic and class subordination, they cling to their whiteness as a guarantee of social status. Thus even the supposedly simple case of prejudice against blacks is more complicated than a hegemony model of ideology suggests; like the case of black anti-Semitism, it involves competing ideologies among groups that suffer varying degrees and kinds of subordination.

Approaching the study of ideology in terms of hegemony rather than justice creates a sixth problem: the familiar but troublesome concept of "false consciousness." This concept focuses not on the question of what is just but on the relationship between individuals' thought and the (objective) interest of the class to which they belong. A person whose beliefs and preferences are contrary to that interest is said to suffer from false consciousness. Hegemonic conceptions of ideology lead inevitably to notions like false consciousness—whether or not they use that precise terminology—because members of subordinate groups often accept their lot and may even oppose political activity designed to undermine the hegemony of superordinate groups.

Inquiries into false consciousness are problematic for four reasons. First, they presuppose that a class can have a unified, objective interest and that the interests of each of the members of that class are not substantially in conflict with it. In other words, the concept of false consciousness assumes without further investigation a particular state of affairs about the benefits of collective

action: it assumes that each individual member benefits sufficiently from promoting the group's interest. But even assuming that African Americans as a group have an objective interest, there is no reason to think that the individual interests of some African Americans might not conflict with that group interest. Indeed, by taking contrary positions they may realize considerably greater benefits personally than they would have if they had adhered to the "party line." It is hard to argue that such persons suffer from false consciousness—indeed, they may see what is in their interests more clearly than many other people.

Second, the concept of false consciousness tends to elide distinctions between long- and short-term interests, in part because it is premised on an underlying historical narrative of eventual liberation. But if one no longer believes in such a narrative—for example, the Marxist narrative of the inevitability of proletarian revolution—the multiple and conflicting interests of persons and groups reassert themselves forcefully. It becomes more difficult to state conclusively that a particular perspective is false consciousness. Rather, people may disagree simply because they balance long- and short-term interests, or group interests and individual interests, differently.

Third, accusations of false consciousness are normally directed at members of subordinated groups that dissent from the analyst's view of what is in their class's interest. But the same logic applied to superordinate groups leads to a paradoxical result: members of superordinate groups that support the dismantling of unjust hierarchies also suffer from false consciousness because they are working against their class's interests in maintaining hegemony. If women who oppose gender equality suffer from false consciousness, so too do men who support gender equality. This paradox arises from the fact that the notion of false consciousness is concerned not with the justice of a position but its relation to the interests of a class.

Fourth, the notion of false consciousness is problematic because it is a holdover from the bivalent oppressor-oppressed model of hegemony that I have just criticized. This model makes little sense in a world in which people have multiple and cross-cutting identities. Even assuming that African Americans and women have objective interests as a class, surely these interests can sometimes conflict. When they do, how can an African-American woman avoid a charge of false consciousness, regardless of the position she takes?

Indeed, accusations of false consciousness are often attempts by one portion of a social group to assert a unitary and objective interest that disadvantages or ignores the claims of another portion or subgroup. Working-class women may be accused of false consciousness by middle-class women when in fact their interests differ because of their class position. Similarly, the interests of African-American women may diverge in important respects from those of white women. Once internal divisions and cross-cutting identities are recognized, the notion of false consciousness threatens to become incoherent or at best self-serving.

The approach that I take in this book rejects the notion of grounding an analysis of ideology in the objective and unified interests of social groups. It asks instead whether cultural software tends to produce or sustain unjust effects. This does not eliminate inquiries into the interests of social groups. But it mediates them through the larger question of what is just for all concerned. Because our primary concern is justice, the notion of false consciousness becomes superfluous. An African American who takes positions that undermine the achievement of racial justice may be acting in his or her personal interests at the expense of the interests of other African Americans; but the important question is whether taking those positions promotes or hinders justice. Moreover, a focus on justice as opposed to objective group interest puts the conflicting claims of social groups in proper perspective, for justice does not consist in each group achieving its interests; it involves accommodating the just interests of all.

Ideological Analysis and Normative Commitment

The seventh and final reason to prefer a definition of ideology based on the question of justice rather than on the question of domination is that ideological analysis is essentially and ineluctably normative and interpretive. A definition of ideology in terms of "domination" tends to disguise the normative commitments of ideological analysis. What constitutes domination cannot be articulated in a purely factual way; it requires a view about what is just and unjust in a society. Moreover, the very concept of domination that one might use to distinguish the ideological from the nonideological is itself an object of ideological disputation.

Consider Thompson's definition of domination in terms of "systematic asymmetry" in power and access to social resources. Although this definition seems to rest on facts about society, it must also rest on a conception of justice. The concept of domination must also include a normative judgment about just and unjust treatment if it is to be of any use in a theory of ideology.

In fact, Thompson's definition would be seriously overinclusive if it rested only on the existence of systematic asymmetries in power between groups. Not all examples of systematic asymmetry in power relations involve unjust domination, and not all beliefs that justify or sustain systematic asymmetries between groups are ideological in a pejorative sense. Some systematic asymmetries between groups are in fact justified. Take, for example, the case of felons. Surely this group is systematically disadvantaged in the United States. Indeed, in the United States, we incarcerate felons and deny them the right to vote.

We would probably not say that say that felons suffer from social domination by the law abiding, although we might contend that particular felons suffer from social domination because they also belong to groups that are unjustly treated. The reason we do not claim that felons as a class suffer from domination is that we believe that the systematic disadvantages these people suffer on account of being felons are fully justified. We are justified in systematically disadvantaging rapists, murderers, and child molesters because they have seriously injured other people. That is why our judgments of social domination necessarily require judgments of just and unjust treatment. What differentiates a dominated group from a systematically disadvantaged but undominated one is the question whether the group's lot is due to some present or previous injustice.

Our judgments about social domination are inextricable from our judgments about justice. People of low intelligence are systematically denied many advantages in the United States, including entrance to elite educational institutions and employment in many high-paying occupations like medicine. We might also note the systematic disadvantages suffered by people who are lazy, disagreeable, shy, unambitious, and untalented. Does the mere fact of these systematic disadvantages mean these groups are also dominated? Not necessarily; it all depends on our theory of justice.

Under some conceptions of distributive justice, one might well conclude that people who are lazy, unintelligent, and untalented are oppressed by the industrious, the clever, and the talented. Suppose, for example, that our theory of distributive justice holds that people do not have rights to the fruit of their talents, and that inequalities produced by the use of these talents unfairly disadvantage those with lesser abilities. Or suppose that we think that purportedly negative qualities like laziness are produced by oppressive social structures and that these qualities would be differently produced, differently understood, and differently distributed if these social structures were altered. Finally, suppose that we believe that negative qualities like laziness are matters of social convention, regularly and opportunistically invoked to benefit certain identifiable social groups. Each of these theories of justice may be controversial in some respects. But they aptly demonstrate that our social judgments about domination and oppression are not judgments about facts but about facts mediated through underlying values. They are complicated appraisals of social meaning with ineluctably normative underpinnings.

Moreover, a systematically disadvantaged group may be unfairly dominated, but its unjust domination may not be coextensive with the full degree of its

systematic disadvantage. Some of the disadvantages its members suffer may be unjust, but others are not. For example, it may be perfectly just to imprison certain types of criminals and to discriminate against them in all sorts of ways, but there is a point at which their punishment becomes oppressive and unjustified. If criminals are denied due process, tortured, or imprisoned under inhumane conditions, they may well suffer from domination or oppression. In addition, if all felons are lumped together in people's judgments, so that less culpable criminals like petty thieves are treated the same as serial killers, this may also lead to injustice toward and oppression of the former subgroup. Mentally retarded persons suffer systemic disadvantages in social power that can be justified to some degree by their limited mental capacity, but some of their disadvantages cannot be justified on these grounds. These disadvantages are oppressive, and ways of thinking that justify such oppressive treatment are the proper concern of a theory of ideology. Here too, we cannot base our definition of ideology on the bare fact of disparate treatment or systematic disadvantage alone. We need a conception of justice to distinguish those parts of a group's unequal treatment that involve unjust domination and oppression from the parts that do not.

In this chapter I have argued that a theory of ideology needs a conception of justice. By this I mean that to understand and describe ideology the analyst must bring to bear her sense of what is just and unjust. However, ideological analysis does not require that the analyst have a full-fledged philosophical theory of justice. Nor does this book offer a complete philosophical account of justice. Most people go through their whole lives without developing such theories, and they are nevertheless able to discuss and reason about questions of justice and injustice. Conversely, well-developed philosophical theories of justice are often too abstract to offer specific judgments about whether particular policies or social conditions are just or unjust.

Finally, as we shall see in the next chapter, the very act of engaging in ideological analysis can change our views about what is just and unjust. We must be open to such changes as a condition of our understanding. So the theory of ideology that I offer in this book is designed to be compatible with a wide variety of different philosophical theories of justice. Indeed, in Chapter 7 I will argue that justice is an indeterminate value that must be articulated through human culture. The many different philosophical theories of justice are but one form of this cultural articulation.

Nevertheless, throughout this book I offer examples that assume that certain positions and social conditions are relatively just or unjust. I do this to clarify my arguments about ideology through concrete examples. But these specific judgments are independent of the theory of ideology I present. And I

would hardly be surprised if my own assumptions about what is just and unjust are not themselves possible subjects of analysis and criticism.

My conclusion that the study of ideology must rely on a conception of justice prefigures the answer to the third of the four questions with which I began this chapter—namely, the interpretive stance that we must take toward the object of our study. And it raises even more urgently the fourth question how to deal with the problem of self-reference, given that ideological analysis can also be applied to the analyst's own thought. These questions form the subject of the next chapter.

Theories of ideology take different normative attitudes toward the object of their study. Generally speaking, these theories fall into two categories, pejorative or neutral. A pejorative conception of ideology sees ideology as necessarily opposed to truth or science. To have an ideology is necessarily to suffer from some distortion of belief, because ideological belief disguises, mystifies, or conceals what is true or what is just. In addition, pejorative conceptions of ideology are usually epistemological, because they oppose ideology to knowledge. A pejorative conception of ideology is sometimes called a "critical" conception, but I use the former term because I wish to reserve the word critical to mean self-referential or self-questioning.1

A neutral conception of ideology, on the other hand, sees ideology as a ubiquitous feature of human thought. Neutral conceptions are historicist or sociological because they connect one's ideology with one's position in a particular culture and history. Neutral conceptions contend that all of us have an ideology of some sort, and that our understanding of the social world is necessarily ideological. The ubiquity of ideology does not mean that ideological thought is distorted or false. As its name implies, a neutral conception of ideology describes ideology in nonpejorative terms. To say that thought is ideological is simply to say that it has certain characteristic features. Truth occurs within ideology, rather than being ineluctably opposed to it.

The distinction between pejorative and neutral conceptions of ideology is sometimes associated with different strands of Marxist thought.² Claims that the proletariat fail to understand their true class interests because they are under the thrall of a dominant ideology employ a pejorative conception of

ideology, as do theories that oppose ideology to science (like Althusser's) or make use of the concept of "false consciousness." In contrast, theories like Lukacs's or Lenin's, which identify ideology with the consciousness appropriate to a class given its position in history, might be thought of as neutral conceptions. In fact, they are neutral more in the sense of being historicist than in the sense of being nonjudgmental.3 Moreover, as Raymond Geuss has pointed out, the neutral conception associated with Lenin and Lukacs actually combines a historicist conception of ideology with a positive conception of a particular ideology—the proletarian revolutionary consciousness.4 At a particular point in history, the proletariat must have a special revolutionary consciousness. (Whether it in fact possesses it is another matter.) This class consciousness allows the proletariat to understand social conditions as they really are and allows it to fulfill its appropriate role in history. According to Lenin, it is necessary for a revolutionary vanguard to instill this consciousness in the proletariat. Moreover, unlike the class consciousness of other groups, the revolutionary consciousness of the proletariat is regarded positively. Note that in a pejorative conception of ideology, there is no need for a special positive conception of ideology because ideology is already opposed to truth or science.

Marxist theories of ideology usually define ideology in terms of membership in an economic class and the objective interests of that class. But theories of ideology do not have to be based on economic class membership or economic class interests. John Thompson's and Clifford Geertz's theories of ideology, for example, are pejorative and neutral, respectively, although neither defines ideology in traditional Marxist class terms. Thompson retains the pejorative perspective that he finds in Marxism but applies it to any form of social domination or exploitation, including domination based on race, class, or gender. Similarly, Geertz identifies ideology with general features of cultural understanding, which are in no way limited to or organized around Marxist conceptions of class membership.⁶

Both the pejorative and the neutral conceptions of ideology have symmetrical advantages and difficulties: each is better at dealing with the problems created by the other. Neutral conceptions of ideology are attractive precisely because they seem nonjudgmental: By noting the existence of different and conflicting ways of understanding the social world and their relationships to people's historical and social situation, neutral conceptions appear to embrace the detached objectivity of the social scientist or the fairness and openness of liberal inquiry. The great advantage of pejorative conceptions, on the other hand, is that they are more compatible with the reasons why people have traditionally been interested in developing a theory of ideology: a concern with how people are led to believe in false or unjust things, and how people's ways

of thinking contribute to or sustain injustice. By their own terms, neutral conceptions of ideology prevent the analyst from focusing on these issues, or force her to do so either unwittingly or sub rosa.

In fact, conceptions of ideology that claim to be neutral usually cannot retain their neutrality for very long. A scrupulously neutral approach to conflicting ideologies would make it impossible for the analyst to pronounce one as accurate and another as distorted. The analyst could not even report that each side has grasped part of the truth, for this would mean that the view of the opposite side is to that extent false and distorted. She would simply have to report that the two ways of understanding social reality disagree and to describe the terms of their disagreement. Even then, it may often prove extremely difficult to articulate the nature of this disagreement in a neutral fashion—that is, without ascribing truth or falsity to one side or the other—because of the interpretive character of judgments about social conditions.

Furthermore, a perfectly neutral conception would make it impossible for the analyst to explain how particular beliefs lead to oppression or injustice, for oppression and injustice are themselves contested terms between competing ideologies. Judgments about what is unjust and oppressive (and to whom) look very different from the perspective of different conceptions of social reality. Indeed, these are the very sorts of questions about which competing ideologies disagree most heatedly. The question of whether and to what extent blacks in America are treated unjustly, for example, looks very different to members of the Nation of Islam and the Ku Klux Klan.

This places the neutral theorist of ideology in a difficult position. Describing the effects of competing ideologies becomes virtually impossible if true neutrality is to be retained. Karl Mannheim, for example, attempted to show that competing ideologies had comparative advantages and disadvantages. Mannheim argued that traditional conservatives could see things about social reality that liberals could not understand as easily, and vice versa. But this approach assumes a perspective from which things are understood correctly and one from which they are understood incorrectly, and this leads us back to a distinction between truth and ideology that is characteristic of the pejorative conception. In the same way, a scrupulously neutral conception makes it difficult to articulate how particular ways of thinking sustain unjust power or are self-serving. These descriptions implicitly rely on conceptions of what is socially real and what is just, conceptions that cannot be neutral with regard to competing ways of understanding the social world.

Thus, although Mannheim's broadest conception of ideology, which he calls a total conception, begins as a nonevaluative study of the forms of thought of a given age, it quickly becomes evaluative and normative. Mannheim recognized this fact explicitly: The "diagnosi[s] of [the thought of] an epoch," he

argued, "though it may begin non-evaluatively, will not long remain so." The ideological analyst will "be forced eventually to assume an evaluative position" because of the evaluative nature of historical understanding: "History is unintelligible unless certain of its aspects are emphasized in contrast to others." A fortiori, if one hopes to understand historical phenomena like oppression or domination, evaluative judgments become unavoidable.

Pejorative theories of ideology do not share these difficulties. They permit (and even require) the analyst to argue that the ideological beliefs of others are false, distorted, or self-serving, or that they lead to injustice or oppression. These theories distinguish between ideological understandings of social conditions and the truth about social conditions; they happily offer normative judgments about the thought of others.

While neutral conceptions have difficulty expressing themselves without reference to concepts like truth or justice, pejorative conceptions generally founder on the problem of self-reference. The problem arises as soon as the tools of ideological analysis are applied to the analyst's own thought. If the beliefs of others are affected by their historical and social position, their apparatus of cultural understanding, and their psychological needs to reduce cognitive dissonance, the same is likely to be true of the beliefs of the analyst. The social and causal explanations of belief formation that the analyst applies to others are no less relevant to the analyst's own mental processes. The relation between ideological analyst and analysand is symmetrical; every ideological analyst can be an analysand to someone else.

The phenomenon of self-reference leads to various versions of what has come to be called Mannheim's paradox: if all discourse is ideological, how is it possible to have anything other than an ideological discourse on ideology? The problem arises because, unlike the neutral conception, the pejorative conception defines ideology in terms of falsity or distortion. If ideology is false or distorted belief, the analyst's understanding of the beliefs of others and the nature of social conditions will be warped and limited by her own ideological thinking. She may view social conditions in a self-serving way, for example, and conclude that people who see things differently labor under ideological delusion.

At first glance, Mannheim's paradox seems irrelevant to the pejorative conception of ideology because this conception denies that all thought about ideology is ideological. Discourse about ideology can be nonideological if it is scientific or true. Some thought accurately grasps what is going on in society, and hence accurately comprehends the distorting character of the ideological thought of others. When an analyst is not laboring under the influence of ideology, her analysis of the ideology of others is not distorted and hence is reliable.

Unfortunately, the distinction between truth and ideology does not solve the problem of self-reference but merely restates it in another way. Our judgments about what is true and what is assigned to the realm of the ideological are no less subject to ideological analysis than any other set of judgments. The analyst's judgments about what is a true account of social conditions and what is ideological distortion may also be distorted or self-serving. The boundaries that separate ideology from truth are themselves an object of ideological disputation.

Pejorative conceptions usually assume a unidirectional model of ideological analysis: the ideology-free analyst locates and criticizes ideology in the ideologically deluded analysand. Disagreements between analyst and analysand about social reality are explained as ideological delusion on the part of the analysand. As Terry Eagleton puts it, under this approach, ideology is like halitosis—it is what the other fellow has.¹⁰ But this unidirectional model cannot be sustained, for as Mannheim recognized, the relationship between analyst and analysand is fully symmetrical.¹¹ The pejorative conception of ideology thus becomes a two-edged sword, which threatens to undermine the analyst's views as well as those of the analysand.

When we dissolve the study of ideology into the study of cultural software, these questions and these problems still remain. Is our conception of cultural software neutral or pejorative, and how does it hope to resolve the difficulties associated with either approach? In fact, the theory of cultural software is based on a third conception, which endeavors to combine the advantages of the neutral and the pejorative conceptions without their disadvantages. This is an ambivalent conception. An ambivalent conception of cultural software views cultural software as simultaneously empowering, useful, and adaptive on the one hand, and disempowering, distorting, and maladaptive on the other. We are ambivalent about our cultural software because we see both its good and its bad points, and we see how these arise from the same sources. An ambivalent conception of cultural software differs from a neutral conception because it does not attempt to be neutral or nonjudgmental with regard to competing ways of understanding the social world; it differs from a pejorative conception because it does not see historically generated tools of understanding as uniformly bad or maladaptive in the sense of promoting injustice. Rather, it views our cultural software as both empowering and distorting, as both enabling and hindering justice.

The ambivalent conception of ideology flows from our earlier discussion of how cultural software is produced through cultural evolution. The tools of understanding are produced through recursion and bricolage; they are cumulative and jerry-built. They are never perfectly designed for the understanding the social world or the many kinds of problems that human beings face,

although they may be good enough for the purpose at hand. The same mixture of advantage and disadvantage occurs when we consider the consequences of our understanding for social justice. The adequacy of our tools of understanding with respect to the promotion of justice depends upon the context in which they are employed; a tool that is more appropriate in one context may be less useful or wildly inadequate in another. Conversely, a way of understanding the social world that is completely misguided as a general strategy (and therefore may tend to promote injustice when so used) may be quite helpful and appropriate in dealing with specific features of the cultural world.

Mannheim's insight about the advantages and disadvantages of contrasting modes of thought anticipates the ambivalent conception. Mannheim argued that even ways of thinking that largely limit our imaginations may be helpful to understand some features of social conditions; what narrows our vision may sometimes also sharpen it.¹² At the same time, this narrowing of understanding proves unhelpful and distorting if we apply it indiscriminately to other features of social life; it may lead us to misunderstand or overlook important features of social conditions. When Mannheim spoke of the adequacy or inadequacy of thought, he did not specifically have in mind the question of justice; he seemed to mean some combination of serving the interests of a particular group and being appropriate to the historical development of society viewed as a totality. By contrast, the kind of adequacy I am concerned with is the adequacy of our thoughts and actions specifically in promoting justice and avoiding injustice. This distinction is important, for ways of thinking about the social world that are helpful in assisting a particular group to gain economic or political power may nevertheless foster or sustain injustice.

Ambivalence is the appropriate attitude to take toward cultural software because it is the appropriate attitude to take toward culture and cultural understanding generally. The tools of understanding are the preconditions of understanding the social world. Yet they also are sources of misunderstanding. Hence the study of cultural software is the study of the curious and unexpected linkages between benefit and disadvantage, empowerment and distortion. It is the study of how the tools of understanding simultaneously create conditions of freedom and domination.

How does an ambivalent conception of ideology deal with the problem of self-reference, or Mannheim's paradox? It accepts the inevitability of self-reference but argues—consistent with the general conception of ambivalence—that this feature of our thought does not necessarily make ideological analysis futile or unhelpful. Quite the contrary: the ability of thought to turn upon itself is a prerequisite for an adequate analysis of ideological thinking.

The problem of self-reference is unavoidable in ideological analysis because this analysis must always be performed by somebody or someone. It must be performed by a subject constituted by certain tools of understanding and not others. Ideological analysis always occurs within the forms of cultural understanding, not outside of them. Cultural software is necessary for the analyst to understand the social world, the thought of others, and her own thought. Thus the analyst's cultural software is not an impediment to her understanding; it is a precondition of her understanding.¹³

In this way, Mannheim's paradox is transformed when it is stated in terms of cultural software. The theory of cultural software accepts—indeed insists—that all discourse about cultural software must involve the use of cultural software, that all thought about the tools of understanding must employ the tools of understanding. Not all such thought is limiting or distorting for the purpose at hand, however, and not all limitations or distortions are relevant in all contexts of judgment. The possibility of self-reference does not raise an insurmountable obstacle to ideological analysis, because the tools of understanding are empowering as well as limiting, enabling as well as distorting. They are not simply the enemies of comprehension but also the conditions of its possibility. They are not merely hindrances to autonomy and self-understanding but also make autonomy and self-understanding possible. Their dual role forms the essence of the ambivalent conception.

Mannheim's paradox is thus no paradox at all; rather, it explicates the conditions under which ideological analysis must necessarily proceed: the tools of social understanding must be used in order to understand social understanding. Self-reference is not a difficulty that must be neutralized or avoided in order to sustain a study of ideology. It is not an exceptional or subsidiary feature of this study. Rather, it is the central predicament of ideological analysis. Like the story of the tongs mentioned in Pirke Avot, the analysis of cultural software can proceed only through the use of cultural software. The study of cultural software is not unavoidably self-referential, it is fundamentally self-referential.

When we employ the tools of our understanding to think about our own tools of understanding, our thought becomes reflexive and recursive. Human thought is thinking about itself, considering the conditions of its own possibility, and the forms and limits of its own adequacy. A subject constituted by cultural software is thinking about the cultural software that constitutes her. It is important to recognize that this recursion in and of itself involves no contradiction, anomaly, or logical difficulty. Nothing in the nature of cultural software prevents us from using it to think about itself. To the contrary, the reflexiveness or self-applicability of cultural software is one of its most significant features. Human understanding—hence human understanding about understanding—is essentially reflexive and self-referential. It can use its own tools to think about its own tools, and equally important, it does use its own

tools to think about its own tools. Our examination of our cultural software is a reflexive study of a phenomenon already reflexive by nature.

Self-Reference and Self-Criticism

This recognition does not make the difficulties of self-reference magically disappear. It may be true that our cultural software is not uniformly distorting or maladaptive. Nevertheless, our conception of cultural software is ambivalent, not uniformly positive. If the tools that we employ to understand social reality are heuristic and have unexpected side effects and limitations, our own understanding of cultural software—either our own or that of others—may be affected by these features. Our understanding and our analysis may turn out to be unacceptably partial, counterproductive, misleading, or unhelpful. Moreover, the positions of the ideological analyst and the analysand are still symmetrical. We may still question the analyst's understanding using the same tools she applies to the analysand. When we examine the thought of another person, the tools of understanding we employ may, in the relevant context in which we use them, be badly suited for the task and may have ideological effects on our own thought. Thus if we disagree with another person about what is going on in society, we must recognize that this disagreement may not be due wholly to ideological effects on her thought but may also be due to ideological effects on our own.

The symmetry of analyst and analysand means that in an ambivalent conception, the analyst must attempt to examine her own thought along with that of the person she analyzes. Thus, if a pro-choice feminist discovers that a large number of blue-collar women in the United States are opposed to abortion, she must not immediately rush to pronounce the thought of these women as ideologically deluded. Rather, she must, as a part of the process of ideological analysis, consider what she might learn from these women about the social conditions they face. She must consider the insights into social reality that they might have, and reevaluate her own views in light of them. Without such an inquiry, she has no way of knowing whether the disagreement between her and the analysand is due to distortions or limitations in the analysand's thinking or in her own.

This obligation flows directly from an ambivalent conception. This conception postulates that the tools of understanding do not uniformly limit and distort the thought of subjects. If so, this must be true for both analyst and analysand. If the analyst is empowered and enabled by her cultural software, she must consider the possibility that the analysand is also enabled and empowered by hers, although in different ways and perhaps to a different degree.¹⁴

Successful ideological analysis is possible because and to the extent that the analyst's tools of understanding enable her to understand social conditions well enough to perform the analysis. For precisely the same reason, however, it is possible that the analysand has a grasp of social conditions that conflicts with the analyst's but is nevertheless equally adequate or even more valid. The analysand may in fact see something that the analyst does not see as clearly. By considering how the analysand's thought might have elements of truth or justice in it, the analyst can attempt to analyze and modify her own views. By using the beliefs and opinions of others as a partial check on the analyst's own, ideological analysis attempts to improve social understanding not only for the analysand but for the analyst as well.

I call this dialectical approach to the study of ideology or cultural software a critical approach. By *critical* I do not mean the discovery of flaws or defects in the thought of another person but rather a process of self-reflection and self-discovery that is part and parcel of the ideological analysis of the thought of other persons. A critical approach is inevitably a self-critical approach.

Although critical examination must always become critical self-examination, most people find it easier to see ideological effects in others than in themselves. In fact, they may be able to grasp limitations in their own thought only by transferring their observations about the limitations of the thought of others and wondering how analogous effects could occur in their own thinking. Alternatively, they may critically examine their own thought only after they have been criticized or attacked by others. Once we begin the process of critical self-examination, our views of the other, and her limitations, may change correspondingly. Thus critical thought returns to the self, although it begins in the examination of the other. Critical self-examination is not, strictly speaking, introspection but rather a process of comparative examination between the self and others. It looks inward by first looking outward.

A critical approach involves critical self-examination, but it is not for this reason a private or individual practice. It is the result of interaction with others in the world, an interaction that may be agonistic as well as cooperative. We may not reexamine our own beliefs until others put them in question. Because of the fallibility of our own cognitive processes, we must, to a large degree, depend upon others for the impetus to critical self-examination, just as we often rely on others for other kinds of knowledge. Thus critical practice is fallible and dependent on contingent circumstance (for example, who we happen to interact with) rather than a source of certainty.

A recurring problem with traditional conceptions of ideology has been that they are unidirectional. They are "critical" only in the sense of taking a pejorative view of the beliefs of others but not in the sense of being self-critical and self-reflective. As a result, these approaches fail to acknowledge the symmetrical positions of the analyst and the analysand. They project the sources of disagreement between analyst and analysand onto the mental processes of the analysand and locate their cause in distortions in the analysand's thought. A unidirectional approach conceives ideological analysis as a critique of defects in the thought of an Other, who is either despised or pitied for them.

Ironically, by failing to understand the views of another as anything other than a distortion, we fail to understand ourselves. The unidirectional approach is the loss of a double opportunity. In contrast, a critical approach recognizes that ideological analysis is not merely the analysis of defects in an Other, in which the existence of such defects is presumed and preordained; it is an analysis of a disagreement with an Other about the nature and justice of social conditions. The disagreement between analyst and analysand is produced by the juxtaposition of contending understandings. These understandings are produced by the use of different tools of understanding or by the use of similar tools in different contexts of judgment. To understand how the disagreement arises, we must try to trace the source of these beliefs in cultural software. Properly performed, the process of ideological analysis must call the analyst's beliefs into question and place them on the table for analysis and scrutiny—a task that can be performed only by using the analyst's own cultural software.

Ideological analysis asks how a particular disagreement about social conditions between analyst and analysand is produced. The answer to this question is not necessarily that the analysand was completely wrong and the analyst was completely right. Rather, the process of understanding how this disagreement arises may affect the analyst's own beliefs and opinions. It may lead her to a deeper and richer understanding of the social world. Yet this process cannot have salutary effects unless the analyst is open to the possibility that her own views are in need of improvement and that the encounter with the analysand has something to teach her. Thus ideological analysis, properly performed, always "risks understanding." To risk understanding is always to risk changes in one's own cultural software. Thus ideological analysis, rather than a form of power or mastery over the analysand, is also a potential source of power over the analyst.

The critical process is by no means foolproof. The study of the causes of disagreement between ourselves and the analysand is not a royal road to truth or an algorithm for intellectual improvement. Indeed, the process of ideological analysis can produce its own ideological effects. One is the possibility that we will not put our own ways of thinking in question—this is the danger of unidirectional analysis, which projects the source of disagreement wholly onto imagined distortions in the analysand's thought processes, and thus preserves

our own thought from ideological scrutiny. Two other types of ideological effects are equally serious. I call these ideological effects hermeneutic conformation and hermeneutic co-optation.¹⁷

Hermeneutic conformation occurs when we interpret the analysand's views in such a way that we believe that she agrees with us. There is no check on our beliefs because we do not think that there is any serious disagreement. Hermeneutic conformation is the production of a false consensus between ourselves and the analysand.

Hermeneutic co-optation arises when we are too eager to assume that the analysand's beliefs are true or more justified than our own. If we too readily assume that disagreements between the analysand and ourselves are due to inadequacies in our own belief, we may come to believe things that are unjust or untrue. We may be co-opted into believing things that we should not believe; our cultural software may be rewritten by this encounter in ways that produce ideological effects in our thought. Hermeneutic co-optation is the achievement of a consensus about the wrong things. It is a special case of the power that understanding can have over a subject.¹⁸

Although these ideological effects are real dangers, they are a necessary risk. Unless we are willing to reconsider our own beliefs through ideological analysis, we can never achieve a critical approach. In any case, our refusal to engage in this process hardly avoids the possibility of ideological effects on our own thought. Quite the contrary, for as we have noted, such a refusal simply projects all sources of disagreement onto imagined distortions in the thought of the analysand; this projection is itself an ideological effect of our own thought.

Among theorists of ideology, Karl Mannheim was the first to emphasize the failings of an insufficiently self-critical conception of ideology; his sociology of knowledge may be viewed as a critical response to the unidirectional analysis inherent in Marxist theories of ideology. Mannheim claimed that the sociology of knowledge must inevitably proceed to the questioning of the analyst's own beliefs and ways of thinking. He argued that knowledge of society was relational—the product of a relationship between the subject, her experiences and position in society, and the object of her knowledge. It follows that the knowledge of the ideological analyst is no less relational. Hence, Mannheim argued, the analyst must put all beliefs, including her own, into question, and ask how their content is related to the thinker's experience and position in society. ¹⁹

Nevertheless, Mannheim's answer to the problem of self-reference was not fully satisfactory. As John Thompson has pointed out, Mannheim's concept of relational knowledge restates the difficulty rather than resolves it.²⁰ Mannheim's other solution argued that the intellectual class would be able to synthesize the competing perspectives of different social groups and hence would

be able to offer a relatively undistorted view of social reality.²¹ This solution was surely unpersuasive when Mannheim first formulated it, and it seems even less persuasive as time passes. As a class, intellectuals seem to be no less prone to ideological effects in their thought than any other group. If they have any special talent in this regard it seems rather to be a special penchant for developing abstract and high-sounding rationalizations for their beliefs and conduct.

The critical approach to ideological analysis is a helpful step toward dealing with problems of self-reference. We can try to use the beliefs of others as a partial check on our own. When we do this, our goal is not necessarily to reach agreement with others; rather, it is to use the project of explaining disagreement as a means of broadening our understanding of the social world. Nevertheless, this solution is hardly foolproof. It does not make the problem of self-reference go away, because many kinds of ideological effects are still possible. Any self-critical practice, no matter how well intentioned, may still be self-serving and hindered by our ways of thinking.

Indeed, I would argue that any approach to ideological analysis that promises to eliminate the problem of self-reference is probably suspect for that very reason. This problem is inherent in the nature of ideological analysis. The best proof of the ubiquity of the problem is the practice of ideological analysts themselves. Ideological analysis almost always has significant blind spots and ideological effects. The history of Marxism is a classic example. By focusing on questions of social class, Marxist analyses of ideology have often overlooked or deemphasized the importance of race, ethnicity, and gender in explaining social injustices. Even Marxist analyses of class relations have often engaged in wishful thinking about the nature of social conditions, the beliefs and interests of the working class, and the likely course of historical development.

This realization places ideological analysis in the same situation as much of our knowledge about the social world. As with all such knowledge, we learn through interaction with the social world and with others in the world. We learn through a process of trial and error. In the final analysis there is nothing special about ideological analysis—directed either at others or at ourselves—that distinguishes it from many other attempts to understand the social world, the thought of others, or our own thought. It, too, is a process of grappling with the world using the tools that lie to hand. Thus we must accept the fallibility of our knowledge about our mental processes (and those of others) just as we accept the fallibility of other knowledge about the social world. Conversely, we must be willing to accept the possibility that our knowledge can be good enough for the purpose at hand if we are willing to subject it to critical scrutiny.

Perhaps the single greatest mistake that we can make in offering an account of ideological analysis is assuming that this form of inquiry (or the form of knowledge derived from it) rests on a higher plane or uses tools more pure and impartial than other forms of cultural understanding. It does not, it need not, and in any case, it cannot. Ideological critique does not stand above other forms of knowledge creation or acquisition. It is not a master form of knowing. It is not the most important or most sure or most perfect form of thinking. Indeed, there is nothing special about it whatsoever; its most distinctive feature may be its utter ordinariness. It uses the same basic tools of understanding that all other social understanding uses. Ideological analysis, and in particular self-critical analysis, employs the tools of ordinary understanding to think about themselves. It is not pure but reflexive and recursive. Even its recursion is not extraordinary, for the tools of understanding are always developed reflexively and cumulatively.

Here once again we may offer an analogy to computers. When a computer boots up, one of the first programs it runs is a diagnostic—a program that checks the adequacy of its informational capacities. The computer can do this only because the nature of its operations allows such recursion—allows various aspects of the software, firmware, and hardware to act as checks against themselves. Far from being a special sort of program, a diagnostic program is in some sense the most ordinary example of a computer program.

There are perhaps no metaphors more misleading than those we often employ to describe the process of self-reflection and self-criticism. These are metaphors of separation and isolation, removal and ascent: we step back, we distance ourselves, we place ourselves above the fray, we rise above our prejudices, we employ disinterested analysis. Given such descriptions, it is no wonder that people assume that ideological analysis is a higher, purer form of thinking. But these metaphors are seriously misleading. Although the study of ideology is necessarily a self-critical study, it does not involve a special method of distancing ourselves from the tools of understanding in order to reflect upon them critically. That is because our tools of understanding are a precondition to understanding and therefore to any reflection on their own adequacy or inadequacy. We are always using some tools of understanding to evaluate the usefulness of others in particular contexts. Our judgments of adaptability and adequacy are necessarily jerry-built and provisional in the same way that all bricolage is.

A critical approach uses our understanding to study our understanding. It tests the adequacy of our tools by the use of our tools. All that we do or can do in these cases is use some of our tools to understand others, and to fashion new tools of understanding in the process. Yet there is no point at which we abandon the tools of understanding so that we might critically reflect upon all of them. Such an attempt misunderstands what a critical approach entails, and

the effort would be impossible in any case. One does not get outside of one's self to understand oneself. If anything, one gets more inside oneself.

Finally, the practice of self-criticism is not disinterested in the sense of being impartial. It is partial by dint of its very constitution by particular tools of understanding. It seeks to be disinterested not in the sense of neutrality but in the sense of fairness or accuracy; yet its fairness is a fairness judged through the use of the analyst's cultural software, and its accuracy is an accuracy measured through the analyst's tools of understanding. Nor is self-critique uninterested or dispassionate; on the contrary, it is a fully motivated understanding—motivated to improve the subject's tools of understanding.

The upshot of this analysis is not a claim that we are not wrong about our ideological analyses of others or even of ourselves. We are sometimes wrong. The point is that we are also sometimes right, or right enough to effect some improvement in our understanding. And we are right not because we somehow escape our cultural construction but because we put it to good use.

We might contrast this account of ideological self-criticism with Stanley Fish's recent attacks on the concept of critical self-consciousness.²² Fish has argued that the idea of critical self-analysis is both sentimental and conceptually incoherent, because it postulates the existence of critical self-consciousness. Fish argues that critical self-consciousness is impossible because it requires one to get outside the forms of own's own thinking in order to reflect critically on what one thinks. Yet one never gets outside the forms of one's own understanding. One is always already understanding oneself using the forms of thought that one currently possesses.²³

In fact, Fish's argument does not prove that critical self-consciousness is impossible. It simply directs us toward a more careful consideration of what a critical self-consciousness might be. Fish's argument gains rhetorical force precisely from the assumption I have been attacking—the notion that critical self-consciousness is a special form of thought that we must develop specifically for the purpose of ideological analysis. This assumption is linked to the metaphor of stepping outside our accustomed ways of thinking in order to reflect on them, and this metaphor is misleading in turn because it suggests a false notion of a self that exists separate and apart from its forms of understanding.

But when we reflect on our own thought processes or consider the adequacy of our own beliefs, we do not need to stand outside ourselves or abandon our tools of understanding. Fish is quite right that we could not do this even if we wanted to. Rather, using our cultural software, we think about what we are feeling, consider what we believe, question our own motives, and compare our views with those of others. We do all these things with the goal of trying to figure out how we think about the social world and how our thought might be improved.

Critical self-consciousness does not employ any special form of cognition outside of the ordinary tools of everyday thought. It is a kind of thought that we are familiar with in everyday life, one that we employ in our most routine dealings with others. We think about the adequacy of our mental processes and our beliefs all the time. We ask ourselves questions like "Was I being polite?" "Did I understand what she said?" or "Am I upset because I am jealous?" We often criticize ourselves for such bad habits of thought as rushing to judgment or forming misleading first impressions. Introspection and selfcriticism are ubiquitous features of our mental life. They are so common that they even have pathological forms, like obsessive self-doubt and refusal to make judgments. Yet we do not introspect by standing outside of ourselves and reflecting on the thing we stand outside of. Rather, cultural software is reflexive: the tools of understanding are tools of self-understanding. To be sure, we have all sorts of mechanisms that are designed to obfuscate and hinder selfunderstanding-for example, mechanisms of ego defense. But this does not mean that self-referential inquiry is not possible. It simply means that it must take place using the tools available and encountering the forms of ego defense that currently exist.

The attack on critical self-consciousness might be taken even further than Fish's version. It might be read as the claim that we cannot improve our understanding of the social world through any process of critical self-reflection. This claim in turn consists of two different assertions. The first is that we cannot change our ways of thinking through critical self-consciousness because we are trapped inside the ways of thinking we currently have. The second is that the idea of improvement seems to refer to a standard of judgment outside of our own current standards, and this is impossible because we can judge only from our current perspective.

It is not true that we cannot and do not change our ways of thinking by thinking about our own thought. The metaphor of software explains why this is so. Our cultural software is constantly being rewritten. It is rewritten through acts of understanding, which means that (among other things) it is continually rewriting itself. Its reflexive features guarantee that it is always the object of its own manipulation. Moreover, our participation in the economy of cultural software described in Chapter 4 presupposes the continual possibility of changes in our cultural software. We change our minds, and our minds change. We have new experiences, and we experience things anew. Because we exist in history, our selves are part of the flux of change, not merely witnesses to it.

We should not offer too sanguine a view of the process of change in our cultural software. If maturity and growth are possible, so are senescence and corruption. If we can be educated, we can also be manipulated. Moreover, the claim that our cultural software changes over time should not be taken to mean that it changes wholly in accord with our conscious design. Just as people make history, but not as they intend, we also fashion new tools of understanding, but not as we design them. The phrase "critical self-consciousness" may tempt us to assume that the mechanisms of critical self-consciousness are wholly within our conscious control. Yet critical self-consciousness is possible only because a great many of our mental operations remain beyond our deliberate control. We cannot consciously control all aspects of our consciousness because the very elements of control must themselves be preconscious. Paradoxically, then, we might say of critical self-consciousness that it can be critical only if it is not fully self-conscious.

The second critique of the notion of critical self-consciousness questions the possibility of improvement. Notions of improvement or regression must be made by some observer. If the observer is ourselves, we are using the tools of understanding we currently have to think about the difference between our past self and the self we are now. Such a notion of improvement is always internal to the way we currently understand the world; we do not employ a transhistorical perspective to make this judgment. Yet this does not show that change does not happen, that people cannot necessarily understand the existence of this change, or that they cannot make acceptable judgments about it. They will simply understand it given the tools of understanding they currently possess. A person who understands Milton better than she did before can also understand that she understands him better. Conversely, a person who has lost the ability to speak Spanish can also understand that she has lost this ability. In other cases, however, the change in our cultural software may blind us to the nature and extent of change. Indeed, this may be so even in the two cases just mentioned.

The critique of critical self-consciousness is valuable not because it shows that ideological analysis is a hopeless endeavor. Rather, it is valuable because it emphasizes the ordinariness and even the banality of the processes by which we understand ourselves and the social world around us. Ideological analysis seems to be special because it is a kind of knowledge about knowledge. Rather than viewing this reflexivity as special, we should recognize it as commonplace. Ideological analysis is not a master discipline that can promise to regulate or direct our understanding of the social world. Rather, it is a form of knowledge acquisition just like the forms it purports to study and critique. It does not regulate the process of discovery without forming part of that process. Its reflexiveness is proof not of its special nature but its ordinariness. This is perhaps the most salutary conclusion of the critique of critical self-consciousness. Once we recognize that ideological analysis is on the same footing as other kinds of

knowledge acquisition, uses the same tools, and even makes the same kinds of mistakes, we will have a more appropriate attitude toward its shortcomings and its possibilities.

Reason as Cultural Heuristic Developed Through History

The theory of cultural software that we have been developing presupposes a conception of reason. Its basic outlines should by now be familiar: Human beings have an innate biological capacity both to reason and to incorporate and develop tools of reasoning, or what I call cultural software. Nevertheless, much of what we call human reason is a cultural product. It is the development of skills and capacities that allow us to make judgments about (among other things) values and social life. The kind of reason we develop through culture is not merely a formal or instrumental rationality; it is a substantive rationality that enables us to make judgments about what is reasonable and unreasonable.

The faculty of human reason is a historical artifact, developed through a collective and cumulative writing and rewriting of cultural software through history. We might call this part of human reasoning abilities the historical or cultural component of reason. It is the result of processes that are both cooperative and agonistic. Through joint effort and struggle human beings strive to name the good and the bad, the true and the false, and to convince others or otherwise impose their beliefs upon them.

In his historical writings, Kant claimed that humanity develops its rational faculties through struggle, a struggle that ends up being cooperative and cumulative without intentionally being so. Kant's conception anticipates the idea of the cumulative creation of cultural software through conceptual bricolage. One should not confuse this process with Kant's generally optimistic view of history. Many useful and noble ideas may be perverted or completely wiped out in the process of cultural change. As we saw in Chapter 2, the development of human reason is an evolutionary process, which makes use of the ability of human beings to possess, use, develop, and proliferate idea-programs or cultural software. We do not know, however, whether this historical process is ultimately a tragedy or a comedy. All that we can say is that it happens.

We have also noted that human beings are partly constituted by their cultural software. Hence what human beings are doing in the historical process of cultural bricolage is constructing both themselves and reason itself. We construct ourselves because we are composed of cultural software. We construct reason because reason has a cultural and historical component: part of what we call "reason"—and indeed, part of what we call "human"—consists of certain tools of understanding that human beings have collectively created over time from more primitive reasoning abilities.

Finally, we have noted that the human capacity for reason is reflexive; that is to say, it can be turned upon itself to change and develop itself. Hegel's anthropomorphism of reason in the form of Spirit can be understood in this way. We can say, along with him, that reason is a historical product that is continually interrogating itself.

The theory of cultural software proposes that understanding of the social world occurs through tools of understanding. We might call these tools heuristics, or aids to understanding. Such an account, of course, must also be a tool of understanding, and it must also be heuristic. There is nothing self-contradictory about such an explanation, however. A problem arises only if one assumes that heuristics are always or necessarily false, that they only or predominantly hinder understanding rather than serve as aids to understanding. The same might be said of symbol or metaphor. Our understanding of the social world occurs through symbol, metaphor, and figural language. Our account of how this occurs must also be described in symbolic, metaphoric, and figural terms. Yet this poses a problem only if such accounts are misleading for the purposes for which they are employed.

The conception of reason that emerges from the theory of cultural software is a notion of reflexive and recursive reason, where software is applied to its own operations. Human understanding about understanding is essentially self-referential. Self-reference can occur with respect to concepts that apply to themselves (the concepts of metaphor and heuristic, for example) or to theories about the thought of subjects that apply to the theorist who pronounces them.

Consider, for example, the present discussion of cultural software. In order to articulate the claims I am making, I have to use heuristics, metaphors, and figures, not only to convey what I mean to the reader but also to understand and express my views on the subject. These tools of understanding, however, are just like all other tools. They are helpful in some contexts and less helpful in others. They are simultaneously empowering and limiting. Moreover, even articulating and explaining this feature of cultural software must make use of heuristics and metaphors. Consider, for example, the figural nature of the terms empower and limit. To empower is to endow with power or force; to limit is to impose an endpoint or a boundary. Ironically, the same phenomenon occurs if we wish to critique the notion of tools of understanding. We might argue that this is an inadequate metaphor or heuristic to explain what we mean by understanding. But in explaining why the theory is inadequate, we must make use of figure and heuristic to express our dissatisfaction. We might say that the theory fails to "correspond," "match," "portray," "capture," or "express" what is really going on.

Thus there is no point in the process of human understanding when one abandons the tools of understanding in order to describe or critique understanding, to articulate or express how these tools operate or malfunction, advantage us or disadvantage us. One does not, in other words, articulate something that is unmediated by tools of understanding and then articulate its relationship to these tools. Rather, one expresses a relation between something already understood through cultural software and the cultural software that one uses to understand it. This relation is itself expressed, articulated, and understood through tools of understanding. What one always has is understanding—which is to say that what one always has is the employment of cultural software.

Nevertheless, I wish specifically to distance myself from the simple assertion that all thought or all reasoning is "just metaphor" or "just heuristic." The problem with such statements is that they too easily devolve into what Ernest Gellner has called "reason bashing." Such interpretations delight in showing the limitations of reason without considering why such a sorry faculty would be sufficiently capable of recognizing its own limitations. By contrast, the theory of cultural software tries to understand how the complex is made from the simple, how the adequate is manufactured from the inadequate, while nevertheless recognizing the side effects and limitations that such a process of development necessarily comprehends.

The terms *metaphor* and *heuristic* have traditionally been freighted with pejorative connotations, perhaps especially so in the case of metaphor and the figural. Before we announce that human thought is just metaphor or just heuristic, we must first understand how it might be possible for thought to involve *just* metaphors and *just* heuristics—that is, metaphors and heuristics that are apt and appropriate, that enable understanding rather than hindering it.

This brings me to a second difficulty with the simplistic claim that thought is only metaphor or only heuristic. Although such a claim seems radical and even dangerous in its pretensions, in fact it is deeply conventional and mired in the same ways of thinking that it purports fearlessly to reject. The use of the words only or just is especially telling. This suggests that there is some other thing that understanding could involve that, unfortunately, poor human reason fails to match. It implies that there are two kinds of understanding, a good, nonmetaphorical or nonheuristic understanding, and a decrepit, figural and heuristic one. It preserves the possibility of a cultural understanding that involves no symbolic intervention but that is direct, unmediated, unalloyed, and unshaped—an understanding that brings no baggage to the act of conception, that does nothing but receives everything, that experiences things as they are, that simply absorbs what is. It preserves the possibility of a reason that uses no tools or devices, that is not a fashioning and weighing, a judging and making sense—and, because it uses no tools, escapes all limitation. In short, such claims dream of an understanding that is not understanding, of a reason that is not reason. And the great irony of this dream is that it is conducted—from start

to finish—through understanding and through reasoning, which is to say that it is conducted through the symbolic and the heuristic, through metaphors like "direct," "unalloyed," "unmediated," and "unshaped." The dream of a cultural understanding without cultural software is the dream of escaping the conditions of understanding; it is the dream of escaping what understanding is.

We should not say that reason is just heuristic and just metaphor. Rather we should say that when reason operates well it employs just heuristics and just metaphors. In the latter sentence the terms *heuristic* and *metaphor* are themselves heuristics and metaphors that attempt to convey the mechanisms of cultural understanding. One of the most intriguing features of the concepts of heuristic and metaphor is that they simultaneously convey the notion of being adequate and inadequate, of being true and false. A heuristic is an aid to understanding rather than understanding itself; a metaphor is a figural description rather than an accurate one. A heuristic is most helpful when it simplifies, which means that under certain conditions it oversimplifies, fails to take into account all relevant conditions, and therefore misleads. A metaphor is most helpful when it reveals an important quality through an expression of similarity, which means that under certain conditions it will emphasize this similarity to the detriment of important differences and will therefore mislead.

The terms metaphor and heuristic are themselves aids to understanding understanding, figures that illuminate the process of understanding. They are thus simultaneously adequate to this task in some ways and inadequate in others; indeed, this characteristic makes them instances of the very kind of things they purport to articulate. The term heuristic is both good and bad at enabling us to understand the kinds of things that are both good and bad at enabling understanding. The term metaphor is both similar and different to what it is compared to-things that express similarity among what is also different. Indeed, the key concepts of this book—those of tool, software, meme, virus, metaphor, and heuristic-all are examples of themselves, and apply both to themselves and to the ways in which they are used. Hence we might expect that they are both helpful in some situations and harmful in others, enabling understanding in some contexts and unduly limiting understanding in different ones. This realization is part and parcel of an ambivalent conception. The concept of ambivalence in the theory of ideology is not simply a claim about good resting on previous evil, or benefits resting on previous harm, and vice versa. It is also a claim about the simultaneous benefits and problems that arise from the heuristic and adaptive features of understanding. And not surprisingly, the ambivalent conception of cultural software—and indeed, the theory of cultural software itself-applies to itself in this way: it has its own benefits and disadvantages, historically created and linked together.

The theory of ideology that I have been developing in this book is based on the principle of ambivalence. Our tools of understanding are partially adequate and partially inadequate to understanding the world and what is just and unjust within it. Yet this idea presupposes that there are degrees of greater and lesser adequacy. It assumes that our thought can be "good enough" under some conditions, even if, in other settings, it seriously misleads us. Similarly, I have defined ideological effects as those that tend to promote or sustain injustices. This definition presupposes that things can be more or less just.

Nevertheless, our judgments of what is just and unjust themselves depend on our cultural software, which is a result of memetic evolution. As I argued in Chapter 1, human values are articulated and refined through culture. Perhaps the concept of justice, like that of truth, is merely the product of a particular development of cultural software. If so, the theory of cultural software faces three serious problems.

First, the process of critical self-reflection would be not only endless but pointless. For our critical judgments would be the arbitrary product of accidental evolutionary developments. Self-reflection would simply be another version of the continuing struggle of different memes to gain ascendancy in our thought processes.

Second, each culture has its own peculiar memetic development. If the idea of justice is merely a product of memetic evolution, perhaps each culture has its own conception of justice or has no conception of justice at all. Ideological analysis requires that we try to see what is just and unjust in the thought of the analysand. But if analysands do not have the same conception of justice as we do, we may not be able to understand their actions properly, for their

concept of justice may be entirely different from ours. At best we will simply impose our own standards of justice on others who do not share them. And it will come as no surprise that, from our perspective, the views of others are found wanting.

Third, just as there may be no common idea of justice between cultures, there may be no common idea of justice between persons within the same culture. No two people share the same cultural software. If "justice" is merely an evolutionary product—a concatenation of particular memes that we have assimilated in our heads—perhaps we are simply imposing our personal conception of justice on others when we criticize their thoughts and actions. Perhaps there is no idea of justice that applies to all human beings—just individual "justice programs" in conflict with all of the others, trying to spread and take over as many different minds as possible.

In short, if we take the memetic development of culture seriously, perhaps justice is an arbitrary mutation, peculiar to each culture's or even to each individual's memetic evolutionary history. Asserting that our judgments of justice apply to other cultures or other persons merely reflects the power of our own memes over our own imaginations. Of course, we may insist that others should look at justice and injustice the same way we do. But that is only because our memes are attempting to dominate and replace the memes in other people's minds. This domination can occur in many ways—by persuasion, by indoctrination, by physical force or economic conquest—but it is at its basis a struggle of memes for superiority and dominance in the minds of human beings. Eventually, perhaps, all human beings may share a common sense of justice, but it will be only as the result of an effective conquest by certain particularly aggressive and effective memes.

Questions like these are serious problems not only for any theory of ideology but for any conception of human morality and politics. Any theory of ideology and any theory of moral discourse must confront them. I have phrased them in the way they arise for the theory of cultural software, but it is clear that much more is at stake in answering them than the fate of this particular theory.

Although it is possible to imagine that justice is peculiar to each culture's or each person's memetic development, it is impossible to be morally engaged with others given this assumption. I shall argue that ideological analysis, and indeed all moral discourse, must presuppose a transcendent value of justice. Tools of understanding produced by cultures to pursue justice are articulations of this value. Because the conception of what is just is necessarily related to what is true-for example, with what has happened and what is happening in society-moral discourse also presupposes a transcendent value of truth.

Defining Transcendent Values

The word *transcendent* has many meanings. For some people it recalls Plato's theory of Forms in the *Republic*: a determinate and universal norm of Justice by which human institutions can be judged and found wanting. Something is just to the extent that it follows the formula or resembles the Form of Justice. But I reject this view, for I do not think that our values of truth and justice are determinate.

By a transcendent value, I mean:

- 1. A value that can never be perfectly realized and against which all concrete articulations and exemplifications remain imperfect or incomplete. A transcendent value is also a transcendent ideal.
- 2. A value that appears to us as a demand or longing. A transcendent value seems to call out to us to enact it in our culture and institutions. Our sense of justice seems to demand that we correct injustices when we recognize them; our value of truth seems to demand that we correct falsehood.
- 3. A value that is inchoate and indeterminate, which human beings must articulate through culture but which is never fulfilled. Precisely because the demand of a transcendent value is inchoate and indeterminate, it can never be completely satisfied. We attempt to realize and understand a transcendent value through its articulations in culture: these include the positive norms of our culture, our technology, and our institutions. But these articulations are always incomplete and imperfect. Our institutions and theories of justice always fall short of what justice demands. Hence there is an ongoing dialectic between transcendent values and their cultural articulations.
- 4. A value whose existence is presupposed by some essential aspect of human life or some essential human activity. Thus the argument for the existence of a transcendent value is transcendental; the existence of the value must be presupposed given the nature of the activity. Hence we can also speak of transcendent values as "transcendental" values.

Not all human values and ideals are transcendent. Machismo and meekness, for example, do not fall into this category. Many human values and ideals satisfy some of the four conditions listed above but not others. It is by no means clear how many transcendent values there are. But I believe that moral and political discourse requires at least two: truth and justice.¹

For some the very notion of truth as a "value" will seem odd. A sentence in a natural language, they will say, can have the logical value of being true or false, but truth itself is not a value in the same way justice is. Many philosophers hold that a sentence in a natural language is true when what it says bears a certain relation to the physical world or to other beliefs we currently hold. These correspondence and coherence accounts miss the phenomenological di-

mension of truth. Truth appears to us not only as a property of sentences but as a demand for understanding and recognition. Thus when I say that truth is a value, I am not attempting to offer an analytic definition. I mean that human beings have an inexhaustible drive to understand what is the case and what is not in the world around them. It is this value that we experience as a demand.

Transcendental Arguments for Transcendent Values

Transcendent values are similar to what Kant meant by regulative ideals: these values are a necessary precondition to certain forms of thought and certain types of activity. Hence the argument for transcendent values is transcendental. A transcendental argument is a "can't help it" argument; it claims that we cannot avoid presupposing something when we engage in a certain kind of thought or activity that we cannot help thinking or doing.

Transcendent values of truth and justice are necessary preconditions to ideological analysis, but one does not have to engage in ideological analysis. However, ideological analysis is really a special case of the more general activities of moral and political understanding and moral and political discourse. This is hardly surprising; the skeptical argument about justice that I offered above does not merely undermine the project of ideological analysis; it also undermines the possibility of moral and political judgment about other persons and other cultures.

To be sure, moral and political judgment and moral and political discourse are not logically necessary—people can live like hermits and have no contact with each other. But they are practically necessary. As soon as human beings come in contact with each other, live with each other, or affect each other's lives, questions of justice between them necessarily arise.

People often like to say that certain beliefs are "true for them" or "right for them" but not necessarily for others. This is a simple way to avoid controversy and appeal to a sense of fair play and tolerance. Particularly if what one believes is likely to be thought unusual or odd, it is easy enough to deflect anticipated criticism by asserting that what one believes is "true for me" or "right for me" but that one wouldn't dream of insisting that the belief has to be true or right for others. This is especially so in a pluralist society like our own, in which respect for differences of opinion (or at least the appearance thereof) is thought to be a virtue.

But the practical difficulty of "true for me" or "right for me" arises precisely when our actions affect other people and come into conflict with other people's values and goals. Then we have to defend what we are doing, either to those we affect or to someone else. At that point we can no longer treat truth and justice like a pie, from which everyone gets to take away his or her own personal and private share. We must regard truth and justice as something that has claims on others besides ourselves. We must abandon the convenient dodge that what we believe is true and right is true only for us and right only for us and for no one else. Of course, we can continue to insist that all individuals have their own truth and their own justice, and that all we are really asking for is tolerance. But then we must claim that our view of tolerance is one that isn't just "true for us" but should be respected and accepted by others as well.

Transcendent ideals of truth and justice are presupposed in our understanding of encounters between people as encounters between subjects of justice—that is, as the sort of entities that can be treated justly or unjustly. Questions about what is true and what is just necessarily arise whenever people affect each other's lives. They arise when people meet together by design or are thrown together by chance. They arise when people live in a single community or when they encounter each other through travel, conquest, or colonization. They arise when people meet face to face in open dialogue or when they affect each other's lives without meeting, as when a bomber places an explosive on an airplane or a factory owner pollutes a river upstream from people he has never met. As soon as we encounter an Other, justice presses its demand on us, whether we respond to that demand or not.

To be sure, people often try to avoid the mutual recognition of others as people who can be treated justly or unjustly. They may refuse to see the people they affect as subjects of justice. So conquering armies and slaveholders have often believed that their victims were less than human; they pretend that the people they subjugate are like inanimate objects to whom no justice is due. Yet even though conquerors and slaveowners refuse to recognize others as subjects of justice, we cannot understand the meaning—and the inhumanity—of their actions until we recognize these actions as an encounter between such subjects. It is ironic but true that we cannot understand the depth of injustice without an idea of justice.

Nevertheless, because the idea of justice is indefinite and indeterminate, the boundaries that demarcate subjects of justice are always contested and unclear. Today, after hard-fought battles over human equality, most people draw that line to encompass all human beings but no other entities. Perhaps in time we will draw it differently, and the nature of the "we" who draws it will change accordingly. But this potential for change simply reflects the fact that our notions of justice are always imperfect and incomplete. The indefiniteness of the boundaries of the subjects of justice is simply another way of expressing the fact that justice is a transcendent ideal.

We should note, moreover, that animal rights advocates might think it possible to act unjustly toward lions and bears without believing that these animals can act unjustly. They might compare the situation of animals to those of newborn infants and certain mental incompetents, who can be treated unjustly but cannot act unjustly toward others. Put another way, they can claim that an entity can be a subject of justice without being an agent of justice. The boundaries of justice must include both the question of who is a subject and who is an agent of justice. The interesting question of whether subject and agent can be separated in the way the animal rights activist describes is beyond the scope of this book. My point, rather, is that the indeterminacy of the boundaries of justice—the indeterminacy of who is a subject or agent of justice—is part of its transcendent character.

Transcendent ideals of truth and justice are transcendental because they frame the structure of our understanding of human action. We need them to understand the meaning of human action in encounters with others, whether this action is directed at us or at third parties, and whether the encounter is friendly or violent, fair or oppressive. Understanding others in dialogic encounters is a special case of understanding human action generally, and ideological analysis is a special case of understanding a dialogic encounter.

To take an extreme example, suppose that a conquering army finds a defenseless group of women and children huddled and starving in the cold. The army then proceeds to execute them and seize their possessions. In this encounter, there is no dialogue, no reasoned analysis, no self-critical doubt and reconsideration. There is only the brute act of power. How is a transcendent ideal of justice presupposed in this vicious act, which allowed no time for dialogue and involved no attempt at mutual understanding? It is implicated in our subsequent understanding of what has happened.

We cannot understand the meaning of this massacre as a human action except by reference to an ideal of justice that applies to both the victors and the vanguished. Even if none of the victims is alive to tell their story, we cannot understand what their murderers did—as the brutal actions of responsible individuals rather than as the random or determined actions of objects—without reference to a common and transcendent ideal of justice. We do not accuse stones and rocks of injustice when they fall in an avalanche and kill innocent people. We do not accuse lions and bears of injustice when they attack people. They cannot act unjustly because we do not regard their action as being of the same order as human action.² What distinguishes the latter kind of action is precisely the fact that it can be just and unjust, and furthermore, that its meaning cannot be adequately understood except against this fact. Because human action is this kind of action, we must presume an idea of justice as part of our framework for understanding it. Finally, we cannot understand the meaning of the massacre unless we recognize that it happened to subjects of justice—to the sort of entities to whom it is possible to act unjustly. The idea of justice frames

our understanding of the meaning of this encounter in terms of both the nature of the action and the nature of its results. That is why justice is a transcendental as well as a transcendent ideal.

Yet, one might object, why do we need to presuppose a transcendent ideal of justice to understand the meaning of what happened at the massacre? Why can't we simply apply the standards of justice of our own culture? Often we do simply point to the positive norms of our culture to judge others. Yet these norms of justice themselves presuppose a transcendent ideal. And when our views are challenged by those who do not share our culture's norms, we will inevitably be led to reassert this ideal.

Suppose, then, after we have condemned the massacre, that the conquerors could speak to us. "What right have you," they might say, "to apply your standards of what is just and unjust to us? By our own culture's standards, what we did was regrettable but necessary. We had the right to do what we did and so we did it. Your culture's standards, parochial as they are, apply only to you and not to us. They can have no claim on us. You think you have understood what happened. But we think you have completely misunderstood and mischaracterized what we did."

To respond to this argument, we must ascend from the positive norms of our culture to a transcendent norm. We must insist that what the conquerors did was unjust not only by our own standards but by a standard that they, too, should agree to; their failure to agree to it shows that they are mistaken, or wicked, or both.

At the moment we make this claim, we must acknowledge that our own views, and the views of our culture, might actually be limited or parochial in some respects. For we appeal to a transcendent standard that might judge and find both cultures' norms wanting. Nevertheless, our ascent to the transcendent norm allows us to turn the conquerors' argument back on them. For we can say to them: "If standards of justice and truth are internal to each culture, you can have no objection to our characterization of you as war criminals. For just as our standards can have no application to you, your standards can have no application to us. We are as correct in proclaiming your evil in our culture as you are correct in proclaiming your uprightness in yours. But your very assertion that we have misunderstood you undermines this claim. It presupposes common values of truth and justice that we are somehow obligated to recognize. And on that ground we are prepared to argue for your wickedness."

I have used an imaginary dialogue to show how our understanding of human action presupposes transcendent values. This use of the dialogic form was no accident. The rhetorical structure of dialogic encounters reveals the regulative nature of transcendent ideals in a particularly striking way.³ Suppose that we find ourselves in a debate with someone about a question of public policy.

Her views are very different from ours. We attempt to persuade her; failing this, we try to persuade a third party that our views are more reasonable than those of our opponent. Our very attempts to convince the audience and justify our own position require that we appeal to common ideals of justice and truth that are binding on both ourselves and the audience. Moreover, we appeal to these common ideals even if we disagree among ourselves about what those ideals require.

We saw previously how our understanding of injustice presupposes an ideal of justice. A similar phenomenon is at work in dialogic encounters. Even when we accuse our interlocutors of great evils, we make reference to a common value of justice that we claim they have failed to live up to. And their defense, even if unconvincing to us, will appeal to reasons that they insist should persuade us and exculpate them. When we criticize our opponent to a third party, we invoke an ideal of justice that applies not only to ourselves and the audience, but to the person we criticize.

In short, transcendent ideals are presupposed by the rhetorical situation of having to persuade an audience. They seem to spring forth magically from the rhetorical encounter. Like a beautiful mosaic whose pattern emerges from the juxtaposition of diverse stones, the framework of transcendent ideals that undergirds the rhetorical situation emerges through the confrontation between different and conflicting perspectives.4

Moreover, these ideals undergird the rhetorical situation regardless of our private intention to tell the truth or to act justly. People often use arguments to deceive each other and convince each other of things that are unjust. They bully and coerce each other with their words. They take advantage of their audience's lack of information, or its emotional, political, or economic weaknesses. Nevertheless, even when we are being deceitful and trying to persuade the audience to believe what we know is not true, we phrase our appeal in terms of values of truth and justice that we claim are binding on both us and the audience. The ideal of truth frames even our act of lying, for the ideal is presupposed by our decision to lie. Without a notion of truth, the practice of deceit becomes incoherent, just as, without a notion of justice, the practice of injustice makes no sense.

The analysis of ideology is a special case of the dialogic encounter. In ideological analysis, we interact with a person, a text written by that person, or a culture. We try to understand their ways of thinking; through this process we learn something about ourselves and our own judgments. This process creates a virtual dialogue with others, even if they are not physically present. Our ambivalent attitude toward cultural software means that although we criticize the others, we must also allow them to criticize us. Thus ideological analysis is a special case of the more general situation in which we are confronted by

people who disagree with us about what is just, and we must deal with their objections through argument and persuasion. Because ideological analysis is a kind of dialogic encounter, it presupposes the same transcendent ideals.

As we saw in Chapter 6, any ideological analysis we apply to the thought of another could, in theory, be applied to our own thought. Because the position of analyst and analysand is symmetrical, we must assume that neither we nor the analysand has a completely accurate or just view of the situation, and that the cultural software of each is partially adequate and partially inadequate to understanding what is just. We must take an ambivalent attitude not only about the other party's cultural software but also about our own.

Thus our ideological analysis assumes that neither we nor the analysand has a monopoly on what is true or just and that neither of our views offers a perfect, complete account. This already presupposes ideals of truth and justice that apply equally both to ourselves and to the analysand and that are not identical with either of our own views. The ambivalent conception of ideology presupposes common ideals against which both of our views might be found partially inadequate.

We must postulate transcendent norms whenever there is a clash or encounter between the positive norms of different cultures, different groups, or different persons. This encounter can be the virtual dialogue of ideological analysis, the actual dialogue of debate and argument, or the physical encounters of politics, warfare, and economics. It can be a genteel discussion or a violent confrontation. In each case, our encounter with an Other causes the transcendent norm magically to spring to life.⁵

Cultural Relativism and Imperialist Universalism

The idea of a transcendent standard of justice might seem to resemble another, importantly different position: I shall call this position imperialist universalism. This is the view that there are universal concrete standards of justice and human rights that apply to every society, whether pre- or postindustrial, whether secular or religious, and that it is the duty of right-minded people to change the positive norms and institutions of all societies so that they conform with these universal norms of justice and universal human rights. This position is worrisome to many people because they see it as a cover for the imposition of a particular set of standards of justice and a particular set of institutions on all of the peoples of the world, whether or not these standards and institutions are appropriate to all, and even if their imposition will result in considerable misery and human suffering. In other words, people are worried by claims of universal standards of justice because they are worried that some form of cul-

tural or political imperialism, particularly from the industrialized West, will be furthered under the name of universal justice and universal human rights. Similar concerns apply within a single multicultural society like the United States. Ethnic minorities, religious minorities, and women may well be concerned that the attitudes, perspectives, and values of majority groups and of men will be imposed on them in the name of universal standards of justice.

All of these concerns are valid ones, but none of them is an argument against the existence of transcendent values. A transcendent value cannot be identical to the values of any particular group because it is an indeterminate standard against which the positive norms of all groups must be seen as incomplete and imperfect. To identify the positive norms and values of one's own group with transcendent standards of truth and justice is already to misunderstand the nature of such a standard.

In fact, one needs to presuppose transcendent standards to make the case against imperialist universalism. The banner of universal standards of human liberty has often been waved in front of conquering armies, whether military, cultural, or economic. People have often invoked universal standards of reason and justice to promote unjust or inappropriate measures that are unfair to foreigners, women, and minority groups. But one can criticize these usurpations only if one presupposes a transcendent standard of justice. For what can the argument against such actions be other than that they are unjust according to a sense of justice that applies both to oppressor and oppressed? Our very notions of tolerance and respect must be based on values that apply both to ourselves and to other peoples and lands.

To criticize imperialism we must argue that the concrete norms and institutions of the West are not universal standards but only imperfect articulations of justice. Applied thoughtlessly to other cultures, they will produce grave injustices. The argument cannot be simply that the imposition of "universal norms" is unjust because these norms conflict with the norms and institutions of the other culture; it must be because such an imposition offends a sense of justice that transcends both the positive norms of the West and those of the other culture, and against which each might be found wanting. The argument within a multicultural society is similar: We cannot simply argue that it is unjust to impose the values and perspectives of whites and males on women and persons of color because each group has its own values and perspectives. For if the perspectives and values of majorities can have no purchase on minorities, how can it be just to impose the perspectives of minorities on majorities? If the former are improper to apply to the latter, why are the latter proper to apply to the former? Rather, one must acknowledge that the perspectives of each group are partial and incomplete, and that each has the obligation to

understand the possible claims of truth and justice in the other's perspective. Like any dialogic encounter, this demand requires an ascent to transcendent values.

We might try to avoid this conclusion in two ways: First, we might argue that it is unjust from the perspective of the majority's own values to apply its values to minorities. But then we have no response if the majority disagrees and insists that from its perspective it is being entirely fair. The majority is surely as good a judge as we are about what its values require of it. Second, we might argue that the majority should try to understand the perspectives and values of minorities because the majority is oppressive and the minority is oppressed. But this argument already appeals to a common idea of justice that applies to both groups. Moreover, we are surely not arguing that in all things the perspectives of minorities are right and the majority's perspective must yield to it. Minorities can also have biased and parochial views of a situation. Moreover, there may be multiple minority perspectives, some of which are in conflict with one another. The argument must be rather that various minority perspectives have important elements of truth and justice in them that are likely to be overlooked by majority perspectives because of the majority's position, interests, and cognitive framework. In other words, the multicultural situation we have been considering is the familiar one of ideological analysis, and the same arguments apply to it.

If a belief in transcendent values does not require a fixed and determinate standard of justice that applies to all cultures, why does it not collapse into cultural relativism? In fact, there are two forms of cultural relativism, one which I accept and another which I reject. The first claims that different cultures have different norms because they have different histories and have faced different problems, and that if we understood the history and problems of different peoples, many of their norms would not seem so strange to us and would even seem justified. This claim is not inconsistent with a belief in a transcendent value of justice; indeed, it proceeds from such a belief. The transcendent value is indeterminate; it has simply been articulated in different ways in different cultures.

This version of cultural relativism treats the other culture with hermeneutic charity—like a person or a text that has something to teach us. In this sense it takes an attitude not fundamentally different from textual interpretation or ideological analysis. Hermeneutic charity toward the norms of another culture is a necessary admonition against rushing to judgment and may even enlighten our own views. It is not identical, however, to the fantastic claim that if we take enough time to understand another culture, we will agree that whatever that culture considers just will turn out in fact to be just. This conclusion is equally inconsistent with an ambivalent attitude toward cultural software. After

all, our own culture's norms are partially inadequate and unjust, and it is likely that the same is true of other cultures. It is entirely possible that when we learn more about another culture, we may conclude that some of its norms and customs are not justified by its peculiar circumstances and history, just as our own culture's history does not excuse all of its present practices.

Nevertheless, even if we find that a culture's norms are partially unjust, there is still the further question of whether it is right for others (including ourselves) to take active steps to abolish that culture's practices. There may be good reasons to avoid doing so. First, intervention might greatly disrupt the society and cause even greater suffering and misery than is created by the current injustices. Second, our notions of tolerance and respect for political self-determination may counsel against intervening in the affairs of another country or another culture. Once again, none of these claims is inconsistent with the postulation of a transcendent value of justice. Indeed, all of them implicitly make reference to such a value.

Moreover, this version of cultural relativism is compatible with the recognition that some different ways of living may be incommensurable with each other, as long as not all are incommensurable. We may not always be able to say whether one way of life is more just than another. But this does not mean that no ways of living are more just than any other, or that no changes in a culture's norms and mores would make it more just. Even if it is not clear how to compare a constitutional democracy in an industrial age with a close-knit agrarian farming community, it might still be possible to argue that the practice of appropriating lands from the weak and the defenseless in the latter community is unjust and should be changed.

A second version of cultural relativism, which I reject, holds that questions of what is just and unjust exist wholly within a given culture. One can judge a culture's justice only by its own norms, because "justice" is by definition something relative to a system of positive cultural norms. Because of our own cultural upbringing we may not be able to help judging other cultures. But applying our ideas of justice to theirs is a category mistake. We are quite literally talking nonsense when we do this.6

This conception of cultural relativism is inconsistent with the notion of transcendent values. For that reason it is also incoherent. If notions of justice are wholly internal to each culture, then no culture can meaningfully object when another takes it over, seizes its lands, and massacres its inhabitants. It can insist that under its own standards what the other culture is doing is wrong. But this should hardly bother the invaders, since these standards cannot by definition apply to them. If applying the aggressor's standard of justice to the victims is a category mistake, then applying the victim's norms to the aggressors is equally so.

The reason why this position makes no sense is that one can always object, and this objection is meaningful. But the very fact of raising a meaningful objection already places the victim in a dialogic encounter with the aggressor. And as we have seen, the rhetorical structure of this encounter presupposes transcendent ideals.

People may be drawn to this second form of cultural relativism because they feel it is important to respect the values of other cultures and as a perpetual reminder that their own culture does not have all the answers. But ironically, these laudable reasons require the presumption of transcendent ideals. We need a transcendent value of justice to respect another culture's norms as well as to criticize them. To understand why the actions of other cultures make sense, we must already believe in a common and transcendent value of justice that their norms partially and incompletely articulate, as do ours.

We can see a similar difficulty in Jean-François Lyotard's vision of justice as constituted by a multiplicity of different and incommensurable language games. According to Lyotard, discussions of justice are like moves in a language game. Each language game has its own standards of justice, and each grows and extends itself by the development of new moves and new rules created through playing the game. The problem, as Lyotard sees it, is the danger of imperialism—one prescriptive system may attempt to dominate and extinguish the others, or fail to recognize their singularity and their claims to exist and to develop on their own through their own internal logics. Thus, in addition to the multiplicity of justices, Lyotard concludes, we must also have a "justice of multiplicity." We must have a justice, Lyotard argues, that "prescribes the singular justice of each game as it has just been situated." The justice of multiplicity allows innovation within each game but forbids "terror"—which to Lyotard means the attempt by one game to take over and dominate most of the others.

It is a tolerant vision, but as Lyotard's interlocutor, Jean-Loup Thébaud, reminds him, it also undermines the notion of mutually incommensurable language games. For now Lyotard has himself become "the great prescriber," ordering the various games of justice to respect each other's boundaries. If justice is purely internal to each particular game, how can any game have the authority to prescribe the boundaries of the others? Such a game must be engaged in its own form of imperialism, like a superpower preventing wars of conquest between smaller countries by the deployment of its own troops.

To avoid this conclusion, we must acknowledge that some forms of oversight are just and others unjust, and this judgment must be able to be acknowledged within each distinct language game. Yet this means that the games of justice cannot be completely walled off from each other but must be interpenetrating. Discussions about justice between games must be recognizable and coherent moves within each game. So all the games must have something in common with each other; they must all be able to communicate with each other, at least on the subject of justice between them. Lyotard and Thébaud end their discussion in laughter, recognizing the incongruity of their situation. But this incongruity also shows us the inescapability of transcendent ideals.

Transcendent Values and Positive Norms

Not all arguments about what is just explicitly refer to a transcendent standard of justice. People often argue about justice by referring to the positive norms of their culture and its institutions. A culture's positive norms of justice, even where they are quite unjust, constitute its attempt to articulate the transcendent value of justice. These positive norms and institutions always presuppose a transcendent value of justice, because it always remains possible for people to criticize their culture's norms and institutions on the grounds that they are not sufficiently just.

In the Jim Crow South, for example, people might have argued that a black man should not sit next to a white woman in a public bus because it is disrespectful or because it is against the law. These arguments are appeals to the positive norms of that culture and to its laws. Yet it was also possible for southerners to understand that these norms and institutions were unfair and to criticize them, even while living in the same culture.

This ability to recognize and critique the injustices of our own culture is another reason why it makes sense to speak of transcendent values. It is hardly surprising that people are often critical of other cultures and their traditions. What is surprising is that people are able to apply their critical focus to their own culture's norms and institutions. To be sure, it is always possible that these criticisms are misguided in any particular case. But the mere ability to articulate them already presupposes a sense of justice against which the norms and institutions of one's society might be understood as imperfect and inadequate.

One might object that when we criticize Jim Crow laws, we are only offering criticisms based on some other aspect of our culture's positive norms. Enforced segregation on public buses is unjust because it offends our culture's commitment to equality, as evidenced in other customs and other laws. So perhaps criticism of our own culture can proceed without the need to postulate any transcendent value of justice. People simply engage in internal or immanent critique of their own culture and institutions by transferring judgments and norms from one aspect of that culture to another.

But this poses an interesting problem: If our judgments of justice come solely from the positive norms of our culture, how do we know that there is a conflict between segregation on public buses and our commitment to equality? Why do we not see the status quo as a perfect accommodation between competing interests, whose differences are resolved differently in different areas of social life? Thus blacks are permitted social equality with whites in some situations but not in others because the balance of competing considerations is different. The mere recognition of a problem as a problem already takes us beyond a simple application or transfer of cultural norms. Our very sense that there is a conflict between a culture's norms already presupposes a transcendent standard of justice.

Even if we concede that there is a conflict between different positive norms of our culture, both the customs and laws of Jim Crow and our egalitarian commitments are equally part of our culture and institutions. If the culture's norms and institutions are the only standard of justice, what makes one norm or one institution more just than the others? By themselves the culture's norms cannot tell us how to reconcile a conflict between them, for both antagonistic elements are present in the culture. To take another example, suppose current law protects blacks but not homosexuals from job discrimination. If we believe that this treatment is inconsistent, why does it follow that the just solution is to give homosexuals protection from discrimination? Why not remove legal protection for blacks?

One might respond to this difficulty by arguing that our commitment to egalitarianism is a more central feature of our cultural norms than our commitment to racial separation. However, this claim can be understood in two different ways. If something is central because it is more prevalent, we refer only to a positive norm. (For purposes of argument I assume that there are noncontroversial ways of judging prevalence. If there are not, this simply reinforces my point.) Nevertheless, mere prevalence does not guarantee the worth of a cultural practice, unless our only goal is to reinforce the most prevalent positive norms simply because they are more prevalent. Racial inequality may be a central and pervasive feature of a regime of Jim Crow or apartheid, for example, but this does not make it a worthy basis for critique of the few egalitarian norms that might remain elsewhere in the culture. By reinforcing the most prevalent practices of a culture we may reinforce its most deeply unfair elements. On the other hand, by "central" we could mean "more valuable" or "more just." In that case our judgment must refer to a transcendent conception of value or justice that informs our notion of centrality.9

We can resolve these puzzles only if we assume that there is a value of justice that our culture imperfectly responds to. Then we can sensibly say that our culture is more just in some regards than in others, and that all conflicts must be resolved in favor of the more just features. Hence even when we engage

in a so-called internal or immanent critique of our culture's norms and institutions, we must presuppose a transcendent value of justice.

Transcendence of Value Versus Transcendence of Position

The notion of "transcendence" might suggest escaping or getting outside of culture in order to judge it. But this way of talking confuses two very different ideas: the transcendence of position and the transcendence of value. A transcendent position is a perspective that is not limited in any way; it is a sort of "God's-eye view." A transcendent value is a value that cannot fully be realized; it serves as a regulative ideal to our understanding. The notion of transcendence is quite different in the two cases; the first refers to a perfect perspective for understanding, while the second concerns the necessary framework for normative understanding in all perspectives.

Because the word transcendent applies to both, it is tempting to try to connect them. From a transcendent position, perhaps we could fully understand a transcendent value. We could know what was really and completely just if we had perfect information and no impediments to our understanding. But in fact, the idea of a transcendent position is inconsistent with the idea of a transcendent value. A transcendent position makes justice determinate, and a transcendent value cannot be made determinate.

Moreover, the notion of a transcendent position is incoherent. To have a perspective is already to be limited in some way. To have a perspective is to have some apparatus for understanding the world. And to have an apparatus means already to see things in some ways rather than others. To avoid the limitations of each possible apparatus of understanding, one would have to have no apparatus at all. And then one would understand nothing, not everything.

Nevertheless, the idea of a transcendent position has been an attractive one in the history of human thought. The reason is not difficult to understand. Intelligent people understand that disagreements arise from differences in perspective, from lack of information, and from the cognitive limitations of position and interest. They recognize, moreover, that the thought of all human beings is limited in one way or another. Thus it seems natural to think that limitation of perspective is the cause of disagreement and uncertainty about justice. If so, then perhaps by removing the causes we could rid ourselves of the effects. We could put our knowledge about what is just on a sure footing by reference to a perfect perspective or a perfect system for judgment.

These considerations explain the appeal of two very common approaches to justice: ideal observer theories and ideal process theories. Ideal observer theories claim that justice is what an observer under ideal conditions would

find to be just. Ideal process theories argue that correct judgments about justice are the product of what emerges from some ideal process of decisionmaking. Thus justice is the product of ideal decisionmaking conditions or a consensus reached under ideal conditions of dialogue.

Ideal observer theories attempt to avoid speaking in terms of a transcendent ideal of justice by speaking instead in terms of transcendence of position. Ideal process theories try to avoid reference to a transcendent ideal by manufacturing justice out of an ideal procedure. The attempt to avoid the transcendent nature of justice creates problems for each theory. Conversely, to the extent that they are successful articulations of justice, they presuppose a transcendent ideal.

Ideal observer theories face two problems: First, they postulate an observer with ideal characteristics working under ideal conditions. Yet our notion of what makes these characteristics and conditions ideal already presupposes transcendent values of truth and justice. Things are just, not because they are so judged by a person with ideal characteristics; rather, these characteristics are ideal because they help a person understand what is just.

Second, the notion of an ideal observer of justice begins to unravel as soon as we inquire into the characteristics of the observer. All observers have a perspective, but ironically, the ideal observer cannot. To have a perspective is already to be a finite being, with a particular history and a particular set of needs, concerns, and desires. Moreover, having one perspective to some degree precludes having others, because some perspectives are mutually incompatible, or are produced by living different kinds of lives, all of which no single person could lead. Is the ideal observer white or black, pregnant or nonpregnant, untouched by violence or the victim of child abuse? It is clear that she or he can be none of these things, for to be any one of them would already shape and limit her perspective. And this is precisely the problem: to have a perspective is to exist as a finite human being with a particular set of commitments and a particular life history. The ideal observer, on the other hand, can have no gender, no history, and no group identification. The observer can have experienced no defining moments in life, can belong to no political party, can adhere to no ideology or worldview. The observer cannot, in other words, exist as a human being.

The great irony here is that our ability to understand justice stems from our situatedness and our finitude. Our life experiences are the raw materials from which we make sense of the normative demands of life. Without them we cannot understand anything at all. What we most have in common with other human beings is what separates us: our finitude, our inadequacies, and our limitations of perspective. We are able to make normative sense of the world because we exist as individuals with a history, who have experienced things and been changed by them, who have perspectives and cultural software that simultaneously limit and

empower our understanding. This finitude, this historicity, this limitation, is what makes the transcendent appear to us as transcendent—beyond our grasp and full comprehension. Ideal observer approaches fail because they rid observation of its humanity, which is inextricably linked to conditions of human imperfection.

Ideal process theories describe justice as the outcome of an ideally fair process or decision procedure. 10 Like utilitarian and deontological theories of justice, these theories are useful heuristics for articulating our sense of justice. But they never completely capture our sense of justice. Ideal process theories presuppose transcendent ideals, they do not produce them. Whatever procedures we employ cannot justify themselves; they must appeal to ideals of truth and justice in order to convince us that they are fair and likely to produce correct conclusions about what is true and just. This problem is similar to that faced by ideal observer theories: Things are just, not because they are the result of an ideal procedure; rather, the procedure is ideal to the extent that its conditions are fair and it leads to just results.

Procedures cannot be determinative of justice because we can always criticize them in terms of the results they produce: the best criticism of the fairness of a procedure is usually the injustice of the results it produces. So a procedure must at best be considered a way of approximating what is just; it is an articulation of justice, and like all articulations, it will necessarily be imperfect, producing results that are always subject to further criticism.

Ideal dialogic theories are interesting and important versions of ideal process theories. They hold that truth or justice is what people would agree to after a dialogue under ideal conditions.11 Earlier I noted that transcendent ideals of truth and justice seem to emerge naturally out of the structure of dialogic encounters. Hence it is natural to attempt to identify truth and justice with what results from ideal dialogue. One might thus identify truth or justice with the actual consensus of the community in the long run.¹² Or, recognizing that many encounters are unfair and coercive, one might identify truth or justice with the consensus emerging from a dialogic encounter under ideal conditions.

Ideal dialogic theories are distinctive because they emphasize that truth and justice are linked to ongoing processes that involve both discovery and creation. The substance of an agreement about truth and justice is true or just not because it conforms to some preexisting test or criterion but because it is the result of a contingent process that results in agreement. The fact of agreement makes true or just what would not be true or just absent the agreement.¹³

Yet even ideal dialogic theories presuppose a transcendent ideal of truth and justice. Such theories well understand that the brute fact of agreement does not make the substance of the agreement true or just. They must distinguish, in Habermas's terms, between a rationally grounded consensus and a false consensus. The task of the ideal dialogic procedure is to make an agreement between finite human beings of limited perspective impervious to criticism on the grounds that it might be unjust or partial. We can try to solve this difficulty by postulating that the agreement takes place under ideal conditions. Yet as before, transcendent ideals of truth and justice are presupposed in articulating the ideal conditions. What makes these conditions ideal is that they lead the parties to an understanding of what is true or just. Thus the fact that the parties lack important information or suffer from unconscious needs to reduce cognitive dissonance tends to make the results of their deliberation suspect. But the reason why they are suspect is not simply that they deviate from the criteria of ideal dialogue. It is because they are likely to produce a consensus that is not true or just.

Ultimately, however, the problem is not simply that we need better procedures; the problem is that we need better people. The participants in any community are finite beings of finite intelligence whose understandings are shaped and circumscribed by their history. Their perspectives are necessarily limited by the partial inadequacy of their tools of understanding and by their inability completely to take into account situations and consequences beyond their apprehension as well as those which may arise in the future. No dialogue between finite human beings, whose understanding is constituted by the historical development of cultural software, can be an ideal dialogue under ideal conditions. For their perspectives are always limited by the fact that each has a perspective. The history of their discussions is always limited by the fact that each has a history. The only truly ideal dialogue would be one between gods. They would already understand everything, and therefore there would be nothing left to say.¹⁴

I believe that there is a deep connection between being the kinds of finite beings we are—who have absorbed tools of understanding produced through evolutionary bricolage—and our experience of moral and factual truth as transcendent ideals. Although people may have used the fiction of a transcendent position to understand transcendent values, the two notions are actually opposed to one another. It is precisely because transcendence of position is impossible that we experience justice as a transcendent ideal. To have a perspective about human action is already to be imperfect, fashioned from what Kant called "the crooked timber of humanity," imbued with tools of understanding that are the product of history and the object of ambivalence. Yet to have a perspective about human action is already to presuppose the transcendent. The transcendent exists because we are imperfect, because we have a perspective. The transcendent is the frame through which we understand the

normative meaning of human action. The transcendent is the limit that shapes our horizon of moral experience.

The Muse of Justice

As limited and imperfect human beings, we cannot stand outside our cultural practices and our cultural software. And our values must be immanent, in the sense that we can express them only through the tools bequeathed to us by culture. How then can we make sense of the transcendence of human values while recognizing that we always make judgments within culture? How can human values be both transcendent and immanent at one and the same time?

In Chapter 2 I argued that we should think of value as a verb, not a noun, as something we do or feel, not something we have. Human values are inchoate and indeterminate urges or demands that are articulated and refined through culture. A transcendent value is a special kind of human value, a value that can never be fully fulfilled. A transcendent value is an inexhaustible demand.

This way of speaking conflicts with the standard metaphor that we employ to describe evaluation, the metaphor of measurement. We evaluate things by measuring them against our sense of justice, just as we measure the length of a table by laying a ruler against it. Hence we have the familiar metaphors of number (things are more or less just), weight (justice comes from balancing competing considerations), size (the lesser of two evils), and distance (coming closer to or diverging from justice). In the standard conception, values work like scales or rulers, and evaluation is a kind of measurement.

This conception has important metaphorical entailments. The first is a separation between value and the thing valued. Because a value is a standard of measure, it must exist apart from the thing that it measures. One cannot use a ruler to measure itself anymore than one can use a balance to weigh itself.

As a result, the metaphor of measurement also seems to suggest that we must somehow stand outside culture in order to evaluate it. We must use a determinate conception of justice existing outside our existing culture to evaluate it, in the same way that rulers must exist independently of the objects they measure. Thus the twin notions of determinacy of value and separation from the object of evaluation are yoked together under the metaphor of measurement.

The idea of justice as an indeterminate or inchoate urge does not fit this familiar metaphor. Instead of a standard of determinate measurement, the transcendent value of justice is an insatiable urge. Thus we have two contrasting metaphors of the value of justice: justice is like a ruler of determinate length that we use to measure the world, and justice is like an indeterminate demand

that can never be fulfilled despite our best efforts. Each of these is a metaphorical account of human value. Each is helpful in its own way, but neither can be usefully employed in all contexts and circumstances. To understand the phenomenon of transcendence we must recognize the metaphor of measurement as a metaphor, and exchange it for a different figure.¹⁵

The contrasting metaphors of determinate measure and indeterminate demand produce different accounts of why our cultural institutions are imperfect, why there is no example of justice in the world that is perfectly just. There are two ways of expressing this inadequacy. One makes use of the notion of a determinate measure, and the second makes use of the notion of an unfulfilled but indeterminate demand.

Under the first metaphor, a determinate conception of justice exists apart from individual examples of justice and is used to measure them. So we explain the fact that no example of justice is perfectly just by saying that each example is an imperfect representation of a determinate conception of justice. The justice of a law or an institution is a question of the quality of the correspondence between the determinate idea of justice and the concrete example. Thus virtue is a process of good copying, and the virtuous person is a good copyist. One makes an institution just by copying the determinate idea of justice as accurately as possible in all of its details. But because no copy can be perfect, there is no perfect example of justice in the world.

Under the contrasting metaphor, justice is an inchoate yearning that we attempt to articulate through our cultural constructions. To be just we must construct examples of justice using the indeterminate urge for justice as our goad rather than as our guide. This means that the virtuous person is not a good copyist but a good architect. She attempts to satisfy her sense of justice by constructing just institutions. There are many different ways of constructing a just institution, depending upon the situation in which she finds herself and the resources she has available to her. Nevertheless, she responds to an indefinite and indeterminate value that can never be fulfilled. Her constructions cannot exhaust justice's demand. Thus human cultural creations will always fail to be perfectly just, but not because they are defective copies of a determinate standard. Their imperfection arises from the necessary inadequation that must exist between an indeterminate and inexhaustible urge and any concrete and determinate articulation of it. This relationship of inadequacy between culture and value is what we mean by transcendence.¹⁶

Note that unlike the metaphor of measurement, the metaphor of the indeterminate demand does not suggest that we must go outside our culture to evaluate our institutions. Rather, we feel the demand of justice as we construct and reconstruct our institutions using the cultural tools bequeathed to us. To feel the demand of justice we do not have to travel to a place beyond culture; the demand presents itself as a sense of the inadequacy of our tools that we experience as we work with them.

We might understand this idea better through an analogy to the myth of the Muse. In Greek mythology, the artist created works of beauty because of a Muse, who not only inspired but also demanded the creation of the work. But there are two different conceptions of the artist's relationship to the Muse. In the first, the artist is no more than an amanuensis who copies out what the Muse dictates to her. Artists often talk this way out of a sense of modesty: they tell us that they merely wrote down what a higher intelligence created. This version of the myth, however, disguises important features of human artistic creation. Most artists are not obedient copyists; they create only as a result of practice and hard work, and they suffer greatly for their art.

In the other conception, the Muse is a harsh taskmistress who relentlessly drives the artist to create the beautiful, often to the point of madness. The Muse demands enormous sacrifices of the artist but is never satisfied with the results, and so artists live their lives in a sort of perpetual bondage to their Muse. The copyist has the advantage of knowing what the finished product will look like; the servant of the Muse does not have this luxury. The servant must turn her inchoate sense and drive for beauty into a work of art, always with the risk that it will not please the Muse who goads her. In this story, the Muse is the mythological externalization of the human drive to value and create works of value. Thus we might say metaphorically that transcendent human values are like Muses; and that there is a Muse of justice as well as one of truth.

Do All Cultures Have a Concept of Justice?

So far I have assumed that when we discuss questions of justice with another person, both of us are speaking about the same concept. But what happens if we relax this assumption? Does this undermine the idea of a transcendent ideal of justice?

Suppose that we come across a culture that lacks a word for justice. Instead, they have a concept they call dharmatzedek, a term that I borrow from the Sanskrit word for duty (dharma) and the Hebrew word for righteousness (tzedakah). (I combine the two words because I do not want the reader to assume that I refer to the concepts of justice in either Hinduism or Judaism.) According to the views of this hypothetical culture, dharmatzedek is a cosmic order of the universe. Social order is a special case of the cosmic order. Things adhere to dharmatzedek when they reflect the proper order of nature, an order that includes not only human beings but animals, gods, and even inanimate objects.

Can we still say that a transcendent ideal of justice is presupposed in our conversation with members of this culture? This is really the question of whether it is possible for us to have a conversation with them about what is just and not just. The fact that they have no such word does not mean that such a conversation is impossible. We modify our existing cultural software all the time in order to understand what others are saying. For example, I have just introduced the concept of dharmatzedek into the present discussion.

Let us suppose that our communications with this culture lead us to believe that by dharmatzedek they mean the concept of natural order that I have described above. Then we will be able to have a conversation with them, for our notion of justice is a notion of achieving appropriate social order and rectifying inappropriate social order. It overlaps with their concept of dharmatzedek even if it is not identical to it. Their sense of social order will be very different from ours, and they may use very different ways of expressing it and making claims about it. But once we create a theory of what their concept means to them, we can begin the process of understanding how the world makes sense to them. Conversely, from their point of view, we will have a very strange concept called "justice," which concerns the order of society and the entitlements of individuals but does not concern their relationship to other things in the universe. They will see our concept of justice as a truncated and mangled conception of dharmatzedek, with an exaggerated focus on the concerns of individuals. But as soon as they formulate a notion in their own language that makes intelligible what we are talking about, they can begin to see that what we are saying makes sense from our perspective, even if they do not entirely agree with it.

Now suppose further that we discover that this culture is engaging in very inefficient forms of agriculture. Not only that: they refuse to engage in practices that would increase productivity. They believe that these practices would show improper respect for the land and disturb the moral order of the universe. They hold this view even though their forbearance means that many of their people will starve. And when we suggest new methods, they refuse to accept them because our methods are against dharmatzedek.

Note that my very description of their objection shows that we can understand why it makes sense for them to object to agricultural innovation. It is intelligible even if we think it mistaken. Conversely, their conversations with us enable them to recognize that our notion of "justice" is an impoverished version of dharmatzedek; this will allow them to understand why we think that one should adopt the new methods, even though they are convinced that we are quite wrong. Once again, the very fact that we can describe our differences from them means that some sort of mutual understanding is possible, even if it is not a perfect understanding.

Nevertheless, we should not assume from our ability to communicate that the other culture really has the same concept of justice as we do and that their concept of dharmatzedek is parasitic on it. We should not assume that our "justice" corresponds to some sort of "natural kind" and their concept of dharmatzedek does not. But if that is so, doesn't this undermine the notion of a transcendent conception of justice? Does it mean that we must acknowledge a separate, transcendent ideal of dharmatzedek, and so on, for each different conception in each different culture?

Before answering this question, we might raise the stakes even higher. Most people think that Plato offered a theory of justice in his Republic. But the concept of dikaiosunē (justice) in fourth-century B.C. Greece is hardly the same notion as the concept of justice we have today. Today in Western democracies we tend to think of justice in terms of getting what we are entitled to, fulfilling our duties to others, and avoiding injury to them. Our contemporary notion of justice is largely organized around the language of individual rights and focuses on interpersonal injury and benefit. Plato does not have this conception. His major concern in the Republic is how the individual fits into the social order. Whereas many (though by no means all) political theorists in the contemporary West tend to view the state as an instrument for fulfilling individual needs and protecting individual rights, Plato's conception regards social order as primary and the well-being of individuals as secondary. For Plato, dikaiosunē is satisfied when the individual exists in the right relation to his or her society. That is one reason why, in the language of contemporary conceptions of justice, Plato's scheme for an ideal city seems so authoritarian to us.

Nevertheless, Plato is one of the founding voices in the Western discussion about the nature of justice. And he is also identified with the notion of a universal transhistorical conception of justice. Yet if Plato's dikaiosunē is not the same as our contemporary conception of justice, perhaps we face the same difficulties in understanding his conception in the Republic as we face in understanding the meaning of dharmatzedek. How then can we claim that there is a single transcendent conception of justice when each culture seems to have a different conception, including those cultures that form the wellsprings of our own?

We can solve this problem by recognizing that our contemporary rightsbased notion of justice is not a transcendent conception. It is itself an articulation of a transcendent ideal. It has come into being at a certain point in history and will no doubt be replaced by some other normative conception in the future. Plato's society-based notion of dikaiosune is also not a transcendent conception but an articulation that reflects the cultural software of his time. So, too, the natural order-based conception of dharmatzedek represents that culture's articulation of a transcendent value. Each culture (and each person) shares this transcendent value, but each articulates it in a different way.

Throughout this chapter I have called this transcendent value "justice"

because that is the word that comes most easily to me, given my cultural situation, my cultural software. And it is also the easiest way to explain the idea of transcendence to an audience that shares most of my assumptions. But if justice, dikaiosunē, and dharmatzedek are all articulations of some higher transcendent ideal, what is the nature of that ideal? What is the common ground of all three notions? What is the cognitive framework that each of the three fleshes out partially and imperfectly, and that makes it possible for each of us to understand the other?

In a more abstract sense, we are talking about the transcendent ideal of a normative order. In each culture the members have a conception of a normative order that includes a notion of the subjects and agents to whom this normative order is relevant. Our notion of justice, the concept of dharmatzedek, and Plato's dikaiosunë are all ways of expressing the normative order that exists between the relevant subjects and agents. The subjects and agents of each normative order are "individuals," and the normative order concerns their proper relation to each other. By "individuals," however, I mean only the sort of sentient beings who can be subjects or agents, not the full-bodied conception of rights-bearing individuals that we associate with contemporary liberalism. In some cultures, animals and gods can be subjects or agents of the normative order. Moreover, in some cultures, what we call inanimate objects can also participate in the normative order because they are not, strictly speaking, inanimate—they are embodiments of or connected to gods and spirits.

Thus each culture recognizes an idea of a normative order, but each articulates it in a different way. There are as many ways to articulate the idea of a normative order as there are possible ways of articulating the relationships between individuals, society, and the universe. In Plato's dikaiosunē, for example, but not in a modern rights-based conception of justice, the normative order refers to the individual's right relation to the state. Notions of individual entitlement are mediated through the language of this relationship.

It is equally important to recognize that some cultural articulations recognize only some of the possible elements of the normative order as salient. The concept of dharmatzedek, for example, includes our normative relationship to the universe. But the modern conception of justice makes a tripartite distinction between human beings, other living things, and inanimate matter. Justice is a relation between human beings (and possibly some animals). It does not extend to inanimate objects. The contemporary West deemphasizes the possibility of a normative order between ourselves and what our worldview sees as inanimate matter, even though there is currently much interest in protecting the environment. We in the West are more likely to argue for environmental protection because of the ways it will affect future generations of human beings or other living creatures than to claim that we have ethical obligations to rocks,

stones, and bodies of water. Even if we recognized a normative relation between ourselves and inanimate matter, I suspect that we would not call it justice. We would give it some other name, like the "sanctity of nature."

Comparing our contemporary notions of justice with other possible concepts like dharmatzedek or dikaiosunē suggests the many different ways that cultures can articulate the transcendent ideal of a normative order between individuals that we understand as justice. Even though justice is an abstract and indeterminate concept, it already articulates and restricts that ideal, reflecting the concerns and attitudes of our cultural moment. Through culture we divide and distribute our sense of the normative order into multiple values and virtues, of which justice is only one among many. Not all cultures will do this in the same way and so we should not expect that their moral language will share the same distinctions as ours. But because all of our moral discourse presupposes the idea of subjects and agents in a normative order, we can be intelligible to each other even if we do not always agree. Indeed, if we could not understand the speech and actions of others as presupposing a normative order with subjects and agents of some kind, it is likely that we would not even understand them as being rational agents.

Pragmatism and Historicism

My argument that different cultures articulate transcendent values in different ways brings me at last to a final objection—one that should be particularly important in a book whose vision of culture is largely historicist and whose view of knowledge is essentially pragmatist. The objection is that one cannot meaningfully speak of ideals or values that transcend cultures because the idea of transcendence is itself wholly peculiar to a particular cultural tradition of discourse—the discourse of Western philosophical thought. The concept of transcendent ideals has a specific history and genealogy that stretches backward from the present day to Kant and to Plato. The idea of transcendence is itself merely one set of tools of understanding that have been developed at a certain point in history to solve particular kinds of problems. Thus, the argument goes, the view of human values as a series of articulations of a transcendent framework neglects the possibility that the very idea of a transcendent conception is itself one of the contingent artifacts of a particular culture. And once we recognize the contingency of the idea of "transcendence"—as a historically produced tool of understanding—we can no longer take seriously the notion that justice and truth are values that transcend all cultural traditions.

The response to this pragmatic objection is entirely pragmatic in spirit. Surely the idea of a transcendent value is a product of a particular cultural history. We can trace its development from Ancient Greece to the contemporary West, and so the particular shape it has taken is contingent in the sense of having been the product of memetic evolution. But it does not follow that the features of the human predicament expressed through this theoretical concept are themselves wholly contingent. Rather, I argue, the concept of "transcendent value" is the best way, given who we are and where we are now, to make sense of these features of human existence—our experience of justice as an inexhaustible demand, and our sense of the inadequacy of all attempts at capturing this value and making it determinate. The concept of transcendence is the most adequate way of describing this inadequacy.

As tools of understanding, all of our ideas are imperfect, and this holds true even for our ideas of perfection. The idea of transcendent values is itself merely an articulation of that which it purports to describe. As an articulation, it is surely subject to revision. And perhaps someday we will exchange the notion of transcendence for another that will be more successful. But this does not mean that the features of human life that our ideas attempt to express are themselves wholly contingent and wholly internal to our discourse. Our conceptions are revisable only because there is something against which we revise them.

Moreover, there is a curious sense in which even the pragmatist objection to transcendence must make use of transcendent ideals. The pragmatist objection is that there cannot really be transcendent ideals because of the historical emergence of the concept of transcendence. Yet such a claim seems to hold itself apart from its own pragmatist scruples. For the objection must surely apply to itself; it is made wholly from within the discourse of a particular culture—and therefore can hardly serve as a judgment about the thought of other cultures. Moreover, the pragmatist objection seems to present itself as an assertion about the way things "really are" that applies with equal force to claims about truth and justice made in other discourses from other cultures with other histories. It offers an impossibility theorem applying to all cultures from within a particular culture. It makes a transcendent claim about the impossibility of transcendence.

The pragmatist thus ends up in a curious reversal. What I have dubbed the pragmatist argument turns out not to be so pragmatic at all, for this argument wants to see behind the illusion of adequacy of a particular conception. It believes in the reality of this illusion and thus in the reality of the state of affairs that the illusion conceals. The pragmatist argument wants to insist that, despite the comfort that the notion of transcendent ideals might give us, they are products of a cultural moment. Hence they cannot describe what is really the case; they cannot truly apply to any other culture than our own.

Conversely, as I have suggested, the argument for transcendent ideals is more truly pragmatic in temperament. Given who we are and where we have come from, the language of transcendence is the best way to explain our ability to discuss questions of truth and justice with other cultures and other persons. It is the best way to understand the phenomenological demands of truth and justice. It is the best way to describe the relation between human values and the felt imperfections of this world. Moreover, transcendent concepts are implicated by many other beliefs about ourselves and our world that we would find it hard to jettison. In other words, the pragmatist argument for transcendent values is that one should accept these concepts and this way of talking because they work.

For my part, this response to the pragmatist objection is as conclusive as I need it to be. It is, as I like to say, good enough for the purpose at hand. I am happy to acknowledge that talk of transcendent ideals of truth and justice is a part of our cultural software that arose at a certain point in history to understand the nature of human action, ideological analysis, and moral discourse. Like other cultural software, it may be revised, sharpened, and even discarded in time. But as of now, I argue that this way of talking is the most adequate way of describing the human predicament. More than that a pragmatist surely cannot demand.

As for the historicist, I would go even further: A historicist conception of human culture and human values not only is consistent with the notion of transcendent ideals, it requires them. By "historicism" I do not mean a theory which holds that the content of substantive values is successively revealed to us through the progress of history. Rather, I refer to the view that people's values are shaped by the historical moment in which they find themselves. Hence as the problems people are faced with change, so, too, do their responses. Historicism in this sense is the temporal counterpart of cultural relativism. The historicist wants us to understand how people in different times and places could have held such radically different views of the world and of human values. She wants us to grasp how it made sense for people to believe in things and hold values that we today find curious or even reprehensible.

Implicit in this project are two assumptions. First, the historicist may wish to present the past as strange and even alien to us, but she cannot present it as utterly unintelligible. Rather, to learn the lessons of historicism, we must seek to understand the past in all of its strangeness and alterity. By definition, an unintelligible past can make no sense to us, and therefore we can learn nothing from it. Indeed, the discovery of an unintelligible past simply leads us to the entirely sensible conclusion that we have not done the work necessary to understand it. For otherwise we cannot know whether the unintelligibility lies in the past or is due to the clumsiness of our efforts to comprehend it. The irony of historicism is that it presupposes the basic intelligibility of the past in order that we may experience its strangeness and difference. Moreover, the

strangeness and difference that it presents is only one step in a larger dialectical maneuver. For historicism also wishes to show us how what we find strange and alien made sense to the people who lived through these times.

This leads to the second assumption implicit in historicism. Morally speaking, the historicist does not want to let us off the hook. She wants to upset our smug assurance that the real reason why the views and values of the past seem ugly and ignorant is that they really are ugly and ignorant. Behind this project is usually a further, deeper agenda: the hope that we can take some critical distance from ourselves, that we can understand that people in future generations will find certain aspects of our practices as strange and abhorrent as we find those of the past. And this agenda in turn harbors two equal if opposite hopes: The first is that the present will come to see that it does not hold all of the answers to questions of value. The second is that if we can learn to be charitable to the strangeness of the past, we may merit an equal charity from the future.

The assumptions of historicism open a virtual dialogue between ourselves and the past, a dialogue that has much in common with the critical approach to ideology that I have offered in this book. To make the past intelligible to us, we must understand why the actions of previous generations made sense to them. We must attempt to see the truth and the justice in what they thought and what they said. And this project brings us inevitably back to the postulation of transcendent values of truth and justice.

Our recognition of historical changes in values requires ideals against which to understand this change. We can describe the history of people grappling with successive tools of understanding that reflect the periods of their emergence because we have a backdrop against which we can describe the limitations of this grappling. We are able to observe the parade of human conceptions passing through history, mutating and reversing themselves, because we have a language for describing their relative adequacy and inadequacy. In this way the concepts of historicism and transcendence are interdependent and intertwined. Each supports the other as its necessary adjunct and companion. We understand the transcendent as transcendent because we can see its articulations vary in history. The variance of history is coherent because we understand it against the background of the transcendent. The experience of historicism makes the concept of transcendence emergent; the concept of transcendence makes the language of historicism coherent.

Our discussion so far has been aimed at dissolving the study of ideology into the larger study of cultural understanding. Now we must take the opposite approach: we must break ideology down into smaller and distinctly analyzable parts. The goal is to replace the study of ideology with the study of ideological mechanisms and ideological effects of cultural software, which, taken together, produce what previous theorists have called ideology.

A great failing of traditional Marxist models of ideology has been that they have usually not offered very detailed explanations of how ideological beliefs are formed in individuals and how these beliefs could be produced by psychological and cognitive mechanisms. Instead, traditional models have tended to focus on the distorting content of beliefs and on whether these distortions tend to serve or disserve the interests of various classes.

By contrast, the theory of cultural software focuses on the mechanisms that produce ideological effects and the means by which they spread widely among human minds. This task requires a study both of social psychological mechanisms and of the ways in which people share their understandings of the social world through the spread of language and symbolic forms.

To this end, the next four chapters offer a partial catalogue of the ideological mechanisms produced by different forms of cultural software. The examples I offer here are not intended to be exhaustive. A comprehensive study of the various devices of human understanding and their possible ideological effects is the work of many lifetimes. Rather, I wish to offer a sample of the wide diversity of tools of human understanding and the many different ways in which these tools can misfire and help produce and sustain injustices.

My goal, however, is not simply to show that different kinds of cultural

software can produce ideological effects. It is also to show that these ideological effects are cultural—that they are widely distributed among members of a culture through memetic transmission. When previous theorists have spoken of ideology, whether in neutral or pejorative terms, they usually have meant shared ways of thinking. For example, a Marxist might hold that many members of the proletariat suffer from the same illusions about the inevitability of their lot; a feminist might point out that most men tend to view certain occupations as inappropriate for women, and so on. But the idea of "shared" ideology brings us back to the metaphysical puzzles that we encountered in Chapter 1. How can these beliefs be shared without presupposing unworkable theories of causation or implausible supraindividual entities? Because traditional theories of ideology have tended to focus on the distorted content of beliefs, or the interests or functions they serve, they have not faced this question squarely.

The theory of cultural software does provide an answer to this question: people experience similar ideological effects because they share similar cultural software and because this cultural software is employed in similar contexts with similar results. If ideological effects are shared, they must be produced by the kinds of memes that can spread widely and reliably through a single population or a group of related populations. This epidemiology produces the effect of shared ideology.

There is no reason to think that memes that produce ideological effects have the same nature or operate in precisely the same way. They need only share an ability to spread widely among human minds. Indeed, the next several chapters will examine very different features of human cognition, including narrative construction, cognitive dissonance reduction, heuristics of decision, metaphor, and metonymy. My examples will be drawn from a wide variety of sources and social scientific models. Often they will have been identified and expounded by theorists with very different theoretical commitments. I shall consider, for example, ideological mechanisms identified by methodological individualists and structuralists, positivists and antipositivists, cognitive theorists and anthropologists, social psychologists and literary critics. It is likely that several of the scholars whose work I draw upon would object strenuously to being discussed together. They would no doubt disagree heatedly about the right way to approach the study of human understanding.

Nevertheless, I must ask the reader to see beyond their various methodological disputes and focus instead on the products of their respective researches. Each of them has, I believe, identified isolated examples of a single, central phenomenon. Each theory, suitably reinterpreted, reveals aspects of human thought with two basic features. First, in each case we have a cognitive mechanism that is "ambivalent": it serves human understanding in some cases and contexts, yet hinders or frustrates it in others. Second, each of these cognitive

mechanisms is cultural, in the sense that it can be and is spread to many different people through communication and social learning. In short, each of these theories reveals a kind of cultural software that, under the right conditions, can act like an ideological virus.

This, I believe, is the proper way to approach such diverse theories about human understanding. The advantage of the theory of cultural software is that it allows us to see how very different research projects can be reinterpreted and united under the umbrella of memetic evolution. Appropriately, this approach is itself a form of bricolage, for it cobbles together different ways of understanding human understanding in the hope of providing a more powerful and unified account. Because human understanding is itself a process of bricolage, we might think of this method as a sort of "metabricolage." 1

The next four chapters consider a wide variety of phenomena, each of which operates differently from the others. The present chapter takes up heuristics of decision and strategies of cognitive dissonance reduction. Chapter 9 concerns the cultural software contained in narratives and scripts, Chapter 10 discusses networks of association, and Chapter 11 explores metaphoric and metonymic models. The variety of these examples amply demonstrates that cultural software does not have to take any particular form to have ideological effects, as long as it can spread widely among a population. Yet as we shall see repeatedly, human language plays an important role. Language is the most effective carrier of memes and is itself one of the most widespread forms of cultural software. Hence it is not surprising that many ideological mechanisms either have their source in features of language or are propagated through language.

Mechanisms Hot and Cold

I begin with heuristics of decision and strategies of cognitive dissonance reduction. Here I shall build largely upon the work of Jon Elster. Elster, in turn, built on the work of two different theories of social psychology: Leon Festinger's theory of cognitive dissonance, and Daniel Kahneman and Amos Tversky's studies on heuristics and cognitive biases.² My goal in this section is twofold. First, I wish to describe the kinds of ideological mechanisms that Elster identified. Second, I want to show the limitations of his approach, and how it can be reinterpreted more broadly and fruitfully in terms of the theory of cultural software.

Elster's work is admirable because he is one of the few theorists of ideology working in the Marxist tradition who has attempted to break down ideology into its component parts and ask how beliefs are caused rather than focusing on the interests they serve. He has turned to the theories of heuristics and cognitive dissonance to provide "micro-foundations" for the Marxist theory of ideology and to put it on a firmer scientific footing.³ Unlike the structuralists I shall discuss in the next chapter, Elster's approach is motivated by a commitment to methodological individualism in the social sciences. Because he hopes to explain as many social phenomena as possible in terms of individuals, their actions, and their beliefs, he tries to explain ideology in terms of individual psychological mechanisms.⁴ His project thus attempts to reduce the Marxist theory of ideology to a theory of individual psychology.

Elster divides ideological mechanisms that distort belief into two groups, which he calls "hot" and "cold." These correspond roughly to the theories of Festinger on the one hand, and Kahneman and Tversky on the other. "Hot" mechanisms are motivational; they are attempts to reduce cognitive dissonance. "Cold" mechanisms are cognitive; they involve heuristics or cognitive biases.⁶

Festinger's theory is only one of several motivational theories that might explain ideological effects. Moreover, the theory of cognitive dissonance has undergone considerable innovation since his original formulation. Cognitive dissonance is now thought to be produced not by the mere fact of conflict or contradiction but by the self's need to preserve its view of itself. Thus, we might call "hot" mechanisms self-preserving or ego defense mechanisms.

In contrast, Kahneman and Tversky's approach explains distorted beliefs not by the need for ego defense but through various failures in cognitive processing. These defects include heuristics or cognitive rules of thumb that operate effectively only in certain limited situations but are extended to situations in which they do not apply. These are classic examples of tools of understanding stretched beyond their usefulness.

Together these two theories of social psychology offer causal explanations for many of the ideological effects described in Marx's writings. For example, one can replace the familiar "dominant ideology" thesis by showing that oppressed groups engage in self-defeating strategies and hold self-defeating beliefs. These strategies and beliefs result from cognitive and motivational biases of the oppressed group rather than being imposed from above by dominant groups. Nevertheless, these biases and illusions may benefit dominant groups, even if they had no hand in producing them.¹⁰

Elster's central example of a motivational bias produced by the need to reduce cognitive dissonance is wishful thinking: people form beliefs because they prefer a world in which the beliefs are true to a world in which the beliefs are false. Elster's formulation of wishful thinking is inadequate because it relies on Festinger's original formulation of dissonance theory. Later theorists have pointed out that mere self-contradiction among one's beliefs does not necessarily produce dissonance reduction because people are able to live with all sorts of contradictions. Only contradictions that threaten the self's view of itself will cause people to reduce dissonance.¹¹ So people engage in wishful thinking

not merely because they prefer a world in which a certain state of affairs is true to one in which it is false. They engage in wishful thinking because accepting that the world is a particular way would significantly threaten their views about themselves. A person might change her beliefs, for example, if holding a certain belief would make her seem less moral, less worthy, less capable, or less in control of her life. People try to preserve belief in states of affairs when they have a personal or existential stake in them.

The need to reduce cognitive dissonance may also cause changes in a subject's values and preferences. One example of this is the phenomenon of "sour grapes."12 People adapt their preferences to value what they believe is potentially available to them. Conversely, they tend to undervalue that which they believe to be impossible or unattainable. Another example of dissonance reduction is the tendency to believe that states of affairs in which we have a stake or which otherwise advantage us are not too immoral or too unjust, or do not show us in a particularly bad light. Lawyers who participate in the adversary system, for example, often have to represent reprehensible clients and argue for positions that they do not believe. They can justify their activities on the grounds that it is demanded of them by the legal system. But if they believed that the legal system was fundamentally unfair, this justification would be seriously undercut and their activities would look much more morally problematic. Hence they have incentives to believe either that the system as a whole is basically fair and just or that precisely because the system is so unfair and unjust to people like their clients, they are entitled to bend the rules to level the playing field and make it more just.

Dissonance reduction seems to explain a number of ideological effects associated with the Marxist theory of ideology. Exploited and oppressed groups, for example, may sometimes believe in the justness, propriety, or adequacy of their fate because this allows them to reduce cognitive dissonance. Such beliefs "may indeed give short-term gratification, but cannot be said to serve the interests of these classes well at all."13 More often, oppressed groups may harbor no such illusions: they may bitterly resent the special favors they feel are granted to more privileged classes. But they may still engage in dissonance reduction. For example, they may openly spurn greater income, more privileged lifestyles, and the symbols of privilege as morally bankrupt or corrupt. At the same time, they may believe that improving their situation is impossible and that their condition therefore must be accepted. Although these distortions in belief may benefit dominant classes, the dominant classes have not caused them. They are caused instead by the need of oppressed classes to reduce the cognitive dissonance produced by coming to terms with the difficulty of their situation.

Elster also argues that Marx's account of the ideological character of reli-

gious belief—that religion is "the sigh of the oppressed creature"—can be explained as a kind of motivational distortion. Hollowing Feuerbach, Marx argues that oppressed classes project the essence of humanity onto a supernatural being; in this way human beings are enslaved to the products of their own imagination. Elster reinterprets this projection as a form of wishful thinking. Human misery causes people to imagine a great and good being who has their interests at heart; this allows them to feel better about their lot because they know that God is watching over them. People create an object corresponding to their wishes and desires; then they understand it as an entity existing external to them so that they can appropriate it through religious devotion or prayer. Again, Elster emphasizes, what is important about this account of religious belief is that it does not explain religious ideology as a function of what serves dominant interests. Rather, it explains religion as the "spontaneous invention of the oppressed, not an ideology imposed by their oppressors." 15

Motivational biases also produce ideological effects in the beliefs of dominant or ruling classes. As Elster points out, one of Marx's most important ideas is that "the bearers of a particular class interest tend to represent it as the general interest of society." Wishful thinking helps explains this phenomenon. The desire to reduce cognitive dissonance causes individuals to believe that what is in their interest is in the interest of society as a whole. It also causes people to have a distorted image of social conditions that support such a view. Thus people can alleviate their sense of guilt or responsibility about poverty by believing that the problems of the poor are exaggerated or that many of their problems are due ultimately to their bad character or immorality.

Finally, wishful thinking can produce strife between classes that might otherwise have common concerns. Middle- and working-class people who face economic insecurity caused by economic restructuring may blame the poor and governmental assistance to the poor for their problems. This both alleviates a sense of guilt or obligation toward those even less fortunate and allows middle-and working-class people to feel morally worthy by comparison.

Cognitive biases explain a different set of ideological effects. For the most part these involve the misplaced use of heuristics; thus we might call them heuristic biases. Like motivational biases, cognitive biases can affect both our views about social conditions and our preferences. A cognitive bias that affects our beliefs about facts is the availability heuristic: "the tendency to believe that the world at large is similar to the part of the world one knows." The availability heuristic assumes that evidence ready to hand is a good source of evidence about parts of the social world not directly experienced by or available to us. 18 A cognitive bias that affects our preferences is a shift in the framing of a problem. When a medical procedure or policy program is described in terms

of its potential gains, for example, it seems more desirable than when it is described in terms of its potential losses.¹⁹

Excessive reliance on the availability heuristic is a particularly common cause of ideological effects. Often it leads to faulty generalization, where "the believer generalizes certain features of his local environment, wrongly believing them to hold in a wider context."20 This phenomenon is reminiscent of Mannheim's point that members of groups tend to extrapolate their experience to all other situations.²¹ A second and related cognitive bias is "conceptual imperialism," which occurs when a thinker uses the categories of her own society "to understand the social structure of other societies or secondary structures within the same society." For example, a thinker might apply specifically capitalist categories to understand precapitalist or noncapitalist social structures. Or she might try to apply American constitutional and political structures to solve the political problems of very different societies with very different histories. Anachronistic thinking and ethnocentrism are familiar examples of this sort of cognitive bias.22

A third form of cognitive bias is the fallacy of composition: "the tendency to believe that causal relations that are valid locally, or ceteris paribus, retain their validity when generalized to a wider context." A special case of this fallacy involves the "natural cognitive tendency to believe that statements which are true from the point of view of any individual agent remain true when applied to the totality of all agents." Because of the fallacy of composition, "there is a natural tendency for the exploited to believe in the inevitability of exploitation."23 The proletariat commit this fallacy when they assume that because they would be worse off without the particular employer who oppresses them, a society without such employers or employment relations would be even worse.²⁴ Similarly, even if a working-class woman in a patriarchal society would be worse off if she were not in a traditional marriage relationship, it does not follow that all women would be worse off if the institution of marriage were significantly changed. The fallacy arises from assuming that the conditions of choice for members of the subordinate group would remain unchanged.

Elster's approach to the Marxist theory of ideology is a genuine advance because it tries to provide causal explanations of ideological effects and because it tries to differentiate ideological phenomena according to their disparate causal sources. Instead of attempting to describe some monolithic entity called ideology, Elster's approach implicitly recognizes that ideological effects result from the confluence of various motivational and cognitive mechanisms.

Nevertheless, Elster's theory is necessarily limited by two features. The first is his adherence to the basic Marxist problematic of economic class.²⁵ In addition to the psychological mechanisms just described, Elster also attempts to

explain the relationship between ideology and modes of economic production. He views the theory of ideology as essentially concerned with economic class.²⁶ This is due in part to his ambition, as the title of his book suggests, to "make sense of Marx." Nevertheless, it is interesting to note that the motivational and cognitive mechanisms that he identifies have no necessary connection to economic class; they work equally for groups defined in ways other than traditional Marxist criteria. Racial and religious groups, for example, may also engage in strategies of dissonance reduction that cause them to engage in self-justifying views of their present situation; ethnic and other social groups may also mistakenly assume that the social world is similar to that with which they are most familiar. Thus the irony of Elster's social psychological approach is that it shows once again how unnecessarily limited is the Marxist approach to ideological explanation.

The second basic limitation on Elster's analysis is that he pays little attention to culture and language as sources of ideological effects. This is partially due to his reliance on social science models that bracket away questions of culture and emphasize individual cognitive processes. It may also be due to his general suspicion of supraindividual entities as explanatory factors. Nevertheless, many ideological effects—including those that Elster is concerned with—occur through the use of shared linguistic and symbolic meanings and shared forms of cultural understanding. Elster's theory of ideology is unduly restricted because he cannot easily assimilate the notion of cultural understanding into his model. In fact, as I shall argue, his own psychological model is actually a special case of what it seems to exclude and marginalize. The heuristic and cognitive biases he describes operate through the use of shared forms of cultural understanding.

Cultural Heuristics for Understanding Human Action

To develop this point, I want to turn first to the work of two quite different social theorists, Paul Ricoeur and Clifford Geertz. Ricoeur adopts a pejorative view of ideology, which assumes that ideology is a distortion, while Geertz adopts a neutral view of ideology as a system of social understanding. What unites these thinkers, however, and what differentiates them from the sort of approach we see in Elster, is that both emphasize that social reality is understood through (or, in their terms, mediated by) shared cultural symbols. Each argues that people use culture to understand their own actions and the actions of others. And each argues that in order to understand how ideology works one must understand that ideology is inextricably linked to the symbolic character of social understanding.

Ricoeur views ideology as a form of distortion of a very special kind. Ide-

ology distorts our understanding of social practices (praxis). Ricoeur argues that people need shared cultural meanings about human action before they can understand their actions and those of others. Hence people always understand social practices and human action through the use of shared symbols. Indeed, social action itself is always mediated through symbolic understandings. Our understanding of what we are doing-an understanding mediated by shared meanings and symbols—is an important component of the choices we make and the actions we perform. Thus, Ricoeur insists, a system of symbolic understandings of social action, or what Ricoeur calls a "symbolic structure of action," must be in place before we can even speak of ideological distortion.²⁷ Shared cultural meanings about human action must already exist before ideology can do its work.

Geertz emphasizes that social understanding—and hence ideological understanding—occurs through the use of interlocking and interrelated figures of speech. The symbolic mediation that is necessary to understanding occurs through tropes similar to the classic rhetorical figures of metaphor, metonymy, synecdoche, oxymoron, and personification. As an example, Geertz offers the statement that the Taft-Hartley Act (a piece of labor legislation opposed by organized labor) is a "slave labor act." A supporter of organized labor who makes this statement is not literally claiming that the Taft-Hartley Act reduces people to a condition of slavery. Rather, she is offering a metaphor.²⁸

Metaphor and other rhetorical tropes are tools for understanding social conditions, describing them to others, and persuading people about them. The persuasive power of a rhetorical figure like metaphor is not diminished by its lack of literal correspondence to the social world. To the contrary, its power arises precisely from the fact that it compares things that are admittedly different and yet the same in some respect. It begins with an obvious difference and ends with a recognition of similarity. A successful metaphor "transforms a false identification . . . into an apt analogy."29

Moreover, the ability of a metaphor to ring true or false depends upon an existing set of cultural associations in which it can be located. Preexisting cultural software determines how arguments are to be framed, how comparisons can be made, and how rhetoric can be wielded. It provides the framework in which the apt description and the inappropriate comparison can occur. Thus, for Geertz, ideology is a cultural system of interrelated associations, symbols, and figures. In other words, ideological effects depend on an individual's participation in a system of cultural meanings and associations. Hence, Geertz argues, "the sociology of knowledge ought to be called the sociology of meaning, for what is socially determined is not the nature of conception but the nature of the vehicles of conception."30

These analyses suggest why Elster's bracketing of cultural and linguistic

sources of ideological effects is too limited. Indeed, what is ironic about Elster's oversight is that he uses metaphors to convey the difference between his two varieties of causal explanation. "Hot" connotes effects produced by emotion or affect, while "cold" implies effects produced absent emotion. Hence we speak of "cool reason" as opposed to the "heat of passion." Elster's own use of these terms is the most telling demonstration that ideological effects can occur as much through metaphor, figure, symbol, and rhetoric—which involve shared categories of meaning—as through individual preservation of the self system or individual cognitive biases or heuristics. These mechanisms cannot easily be fit into his system.

The limitations of Elster's analysis become most apparent in his treatment of tradition and the practice of borrowing traditional symbols in political discourse, a practice that produces "the apparently conservative character of many revolutions." In his discussion of Marx's essay *The Eighteenth Brumaire of Louis Bonaparte*, Elster describes borrowing from tradition as a source of ideological distortion. He attempts to fit it into his model as an example of conceptual imperialism. He argues that it is produced by misuse of the availability heuristic. Nevertheless, Elster admits that this causal explanation of ideological distortion seems different from the others he has advanced because "it offers an explanation of men's conception of the future in terms of the historical tradition, not their present position." The problem in these cases is that people "have to do with the conceptual luggage they carry with them, even at the very moment they grope around for a way to jettison it." "

The ideological distortions created by the use of traditional symbols involve more than idiosyncratic failures of cognitive processing. Elster is describing tools of understanding that involve or employ shared cultural meanings. Thus these tools of understanding are more than individual heuristics—they are shared or *cultural* heuristics. As heuristics, they cannot be wholly detrimental. As noted previously, a heuristic and a bias are simply two sides of the same coin; what is a cognitive heuristic or aid to understanding in one context or situation can also be a cognitive bias or distortion in another. This is simply another version of the argument about conceptual bricolage that underlies the ambivalent conception.

Moreover, as cultural heuristics, these tools are partly constitutive of individuals. To say that people are situated in culture is also to say that cultural tools are situated in them. As we saw in Chapter 1, this puts the concept of "tradition" in a very different light. Tradition is not simply something we live within; it is something that lives within us.

Being part of a cultural tradition is a condition of historical existence. To exist as a historical being is to have a set of tools available to hand that are the legacy of the past. Existence in human history (as opposed to the natural ex-

istence of a mountain or a glacier) is existence in culture. It means that one is composed of cultural heuristics shared by others who are similarly constituted by them.

In the Eighteenth Brumaire, Marx takes a largely negative view of the cultural tools that constitute tradition. Near the beginning of the essay he offers his famous statement that "the tradition of all the dead generations weighs like a nightmare on the brain of the living." Marx thought that "precisely in such epochs of revolutionary crises [people] anxiously conjure up the spirits of the past to their service and borrow from them names, battle slogans and costumes in order to present the new scene of world history in this time-honoured disguise and this borrowed language."33 Although in this particular essay Marx is critical of the use of the past, and of the mawkish and opportunistic use of political props from the past, there is a larger sense in which (as Marx himself emphasized) the phenomenon of borrowing from the past is unavoidable and inevitable. Indeed, political grandstanding—like a French senator dressing in a toga to call up the idea of the Roman Republic, or a modern American politician identifying himself with the Founding Fathers—works only because there is a shared stock of cultural symbols that facilitates a shared political life.

This suggests an important connection between cultural heuristics and the public rhetoric that we use to persuade and influence others. As Aristotle noted long ago, the successful rhetorician builds upon what the rhetorician and the audience have in common.³⁴ And what the two have in common are shared cultural meanings and symbols. Thus rhetoric is also a sort of bricolage, and the skillful orator is a sort of bricoleur: out of the old and familiar, she constructs the new and persuasive. That is why metaphor and figural language in general are effective—because they relate the new to the old, the strange to the familiar, what we come to know to what we already grasp. Indeed, even Marx uses metaphor to make his points: He speaks of the past as weighing "like a nightmare on the brain of the living"—a mixed metaphor to be sure, but a metaphor nevertheless.

Our search for ideological mechanisms thus takes us far beyond the limited confines of Elster's model, important as that model remains. Indeed, our analysis places the heuristics and biases that Elster identifies in a new light. If human action and social practices are always understood symbolically, then Kahneman and Tversky's heuristics and cognitive biases are already embedded in and make use of a set of shared meanings about human action. These heuristics and biases already possess a cultural component; they are already a kind of cultural heuristic.

Elster tried to view the cultural heuristics of traditional thought as a special albeit exceptional case of his psychological model of ideology. In fact, it is quite the other way around. We should rather try to think of social psychological mechanisms, cognitive heuristics, and biases as special cases of a larger category of cultural software that includes many different types of shared cultural meanings and symbols.

Consider, for example, the mechanism of dissonance reduction. The literature on cognitive dissonance has gradually come to recognize that what lies at the root of dissonance reduction is the preservation of the "self system." Mere contradiction or conflict does not lead to dissonance reduction unless the self's view of itself is threatened. But what is the source of the self's view of itself, and what is the source of what the self regards as a bearable or unbearable conflict? In large part the source of both must come from culture and cultural norms internalized by the individual. Put another way, the preservation of the self system involves the self's looking at itself through the eyes of what it imagines others in its culture would think about it. This process is in many ways reminiscent of Mead's idea of a "generalized other" that shapes individual behavior and conscience.³⁵ Thus, although the basic mechanism of dissonance reduction is individual, the content and context of what drives this mechanism is cultural and social.

Moreover, we can describe the mechanism of dissonance reduction in explicitly memetic terms. Each individual mind is a kind of ecology, more hospitable to some memes than to others. Beliefs that do not fit well into the existing ecology of the mind are more likely to be altered, rejected, suppressed, or forgotten. Although the ecology affects the kinds of memes that will survive within it, the memes it absorbs also can affect the ecology itself. Thus new experiences sometimes alter existing beliefs, and new beliefs sometimes are altered to conform with beliefs already held. Strategies of dissonance reduction adjust beliefs and attitudes so that they can survive together in the existing ecology of the individual mind. People whose ecologies are similar—because of their common interests, their common situations, and the commonality of their previously existing beliefs—will provide similar ecologies for new memes. Thus they will tend to engage in similar strategies of dissonance reduction.

We can also interpret "cold" mechanisms consistently with the theory of cultural software. In their research on heuristics and cognitive biases, Kahneman and Tversky do not discuss whether these heuristics or cognitive biases are "hardwired" or whether they are culturally generated. Nevertheless, there are good reasons to think that at least some of them are a kind of cultural software transmitted from person to person. First, individuals can learn to avoid these cognitive biases when they are pointed out. Second, only a certain percentage of individuals fall prey to these errors and biases in psychological experiments. This suggests that these heuristics are learned, adapted, and adopted in new situations and, conversely, that people also can learn when these heuristics are badly adapted to solving particular types of problems.

Of course, neither of these facts conclusively proves that all of these heuristics and biases are produced by transmissible cultural software. To the contrary, some heuristics and biases may produce ideological effects because of the absence of cultural software. Some meme complexes may act like cultural "patches" that allow people to work around the deficiencies of their hardwired heuristics. In that case the reason why only some people fall prey to errors is that they have not yet assimilated the necessary patches through social learning. So we must assume that the group of cognitive heuristics and biases comprehends some combination of the cultural and innate. Nevertheless, even socalled hardwired heuristics and biases depend on shared cultural meanings and concepts of human action to do their work. Transmissible cultural software still may be a necessary condition for most ideological effects to occur, even if it is not a sufficient condition.

Nothing in what I have said suggests an abandonment of Elster's basic thesis: a theory of ideology must attempt to offer causal explanations of ideological effects. My point is that causal explanations cannot bracket away the realm of culture. As Jerome Bruner puts it, "In the end, even the strongest causal explanations of the human condition cannot make plausible sense without being interpreted in light of the symbolic world that constitutes human culture."36

The many ideological effects produced by language and culture fall into the category of cognitive or "cold" mechanisms. They are ideological effects produced by the mediation of social understanding through language, metaphor, narrative, and other symbolic forms. Indeed, they constitute a much larger category of effects than the examples that Elster offers as paradigmatic of "cold" or cognitive mechanisms. Like other cold mechanisms, the tools of linguistic or cultural understanding operate as a kind of heuristic that can produce ideological effects in particular situations. Moreover, it is not easy to separate out their beneficial from their harmful uses. The possibility of ideological effects is built into the very concepts and structure of symbolic understanding and discourse within a culture. Yet once again, while studying these tools for the ideological effects that they produce, we must not forget the extent to which such tools are empowering or enabling. Psychological and cultural heuristics are just that—heuristics that under some conditions perform well enough as rough guides to reasoning but that are misleading in other contexts. The study of ideology, then, might be summarized as the study of "when good heuristics go bad."

To study the many features of human cognition and cultural understanding that can produce ideological effects, we have to cast a wide net. We have to bring together many different fields of study and many diverse types of theories about cultural understanding and the social construction of thought. This

should hardly be surprising, for the theory of ideology is a necessarily and fundamentally interdisciplinary endeavor. In addition to the Marxist tradition, the sociology of knowledge, and the various theories of social psychology just discussed, we might also include structuralism and semiotics, Pierre Bourdieu's theory of habitus, Hans-Georg Gadamer's theory of tradition and horizon fusion, Wittgenstein's theory of language games, Foucault's theory of discourses, and the classical theory of rhetoric, as developed by Aristotle and many others.

Nevertheless, this synthetic project of metabricolage is complicated by two factors. First, the motivations behind these various theories are often considerably different. Some of these theories (like Wittgenstein's or Gadamer's) are philosophical accounts of cultural understanding. They attempt to evaporate certain confusions concerning how understanding works. Nevertheless, they do not offer anything like a causal or evolutionary explanation of ideological effects.

Second, many of these theories (for example, Wittgenstein's and Foucault's) are either unconcerned with or actively hostile to offering accounts of the internal processes of the human mind. Instead, they view features external to the mind—like social behavior or symbolic forms—as the object of their study. This, in my view, is the most serious failing of the theory of discourse that has come to replace the theory of ideology. When discourse is viewed as the structure or content of messages or practices, it casts the study of cultural understanding out of the mind and into the world of behaviors, writings, and articulated symbols. Yet one must do more than identify particular discourses and their structures and effects. It is also necessary to ask how these discourses could be produced by individual human minds, and how what produces them could, in turn, be produced and reproduced in many other individual human minds.

There is a great irony here. The interpretive turn in the human sciences understood itself in part as a rejection of behaviorism. It emphasized the cultural features of human action and the importance of culture and symbols in structuring human behavior. Nevertheless, it is quite possible to re-create a sort of behaviorism within an interpretivist approach if we focus only on behaviors and symbolic forms that are external to the mind, or if we treat the mind as a black box that simply produces and is affected by these symbols and behaviors. The interpretive turn, which we find in thinkers from Geertz to Foucault, has emphasized the role of symbolic forms and culturally meaningful behavior. Yet if these forms are symbolic, they must be symbolic to someone who is able to process and use symbols. If behavior is culturally meaningful, it must be culturally meaningful to a particular person who has some mechanism for making and understanding meaning. To understand the phenomenon of ideology, then, we must marry two separate movements of the second half of

the twentieth century. The first is the interpretive turn, which emphasized the importance in human life of culture and the symbolic. The second is the cognitive revolution, which emphasized the internal processes of human understanding. Each of these movements offers something that the other has either downplayed or disregarded. The theory of ideology—which is only a subset of a general theory of cultural understanding-must make use of each of these approaches and bring them together.

Our theory of how ideological effects are produced requires us to search for cultural heuristics with three basic features: First, they must be transmitted through social learning or communication. Second, they must be stored in the memory of many different individuals. Third, people must use them to reason about the social world. Narratives fit these three criteria particularly well. Narrative and narrative structures are ubiquitous and pervasive features of cultural life. They are easily transmitted through communication. They are deeply entrenched in human thought. Indeed, as we shall see in this chapter, human thought uses narrative structures for a wide variety of purposes.

Narratives are pervasive forms of human thought because narrative structure is a particularly efficient form of human memory storage. Our minds are comparatively well designed to remember and understand narrative sequences. For example, people are better able to recall complex sequences of events in stories than complex lists of words and numbers. Indeed, translating information into narrative form is often an excellent method of memorization. This fact explains the importance of bards and epic poets in oral cultures, where information storage through writing is difficult, costly, or unavailable, and memorization skills are at a premium.

Human beings pick up narrative structures easily from watching and observing events. We naturally seem to create narrative explanations for events or abstract narrative structures from our experiences. We glean narrative structures from life; we impose narrative order on the world. For all of these reasons, the memes associated with narrative structures find a particularly hospitable environment in the ecology of human minds.

Narrative memory is memory of expectations of events in time. It is more

than an ability to recall strings of sequences of events; it also involves the ability to store expectations about what usually happens under certain conditions. These expectations are coded in narrative form. People recall that A happened and then B happened, but they also remember that C is usually followed by D.

Our comparative abilities for narrative memorization have probably been shaped by evolutionary forces. Narrative memory structures are particularly useful for remembering what kinds of things are dangerous or advantageous, making complicated causal judgments about the future, determining what courses of action are helpful, recalling how to do things in a particular order, and learning and following social conventions that require sequential or scriptlike behavior. In the struggle for survival, storing sequences of events and expectations may have proved much more useful than storing isolated bits of information in propositional form.

Whether or not there is an evolutionary advantage to narrative memory, human beings have a particularly well-developed capacity for it. As a result, people use narrative structures for many different mental tasks and operations. These multiple uses are examples of cultural extapation or bricolage—a mental ability or characteristic developed for one purpose is now put to many different purposes. And this particular extapation has far-reaching effects on the development of human culture.

Here are only a few of the things we use narratives for:

- 1. Remembering events in temporal sequences.
- 2. Ordering and organizing the past.
- 3. Explaining human action in terms of plans, goals, and intentions.
- 4. Understanding our own selves and motivations through autobiography.
- 5. Giving causal explanations of events.
- 6. Creating expectations about the future.
- 7. Internalizing expectations about how to behave in social situations and interact with others.
- 8. Providing scripts that tell us how to understand social situations, engage in social conventions, and assume social roles.
- 9. Creating notions of what is ordinary and extraordinary, expected and unexpected, canonical and deviant in social life.
- 10. Accounting for deviations from what is ordinary, expected, or canonical.
- 11. Creating social myths and shared memories that unite groups we are a part of, frame their experience of contemporary events, and produce shared expectations about how the group is supposed to behave.

In short, narrative is simultaneously a method of memory storage, a method of framing and organizing experience, a method for indexing and retrieving information, a method of internalizing cultural expectations, and a method of explaining deviations from cultural expectations. Because narrative is such a ubiquitous tool of understanding, it can also be the source of many different and powerful ideological effects.

Narratives as Networks of Expectations

In general, narrative thought organizes the world into a sequence of events, involving characters and their actions. This is the "plot" of the narrative. The plot and its constituent elements define each other: the plot situates and makes sense of the characters, actions, and events, and these in turn help constitute the plot. Usually the characters in a narrative have reasons for what they do, and their actions have goals. The narrative either assumes or directly ascribes purposes, beliefs, and intentions to the characters. Nevertheless, purely causal stories—for example, the gradual creation of a canyon due to water erosion—are also narratives, although they involve no human characters. Often there are anthropomorphic elements in such stories—we ascribe actions to particular inanimate "characters," like a river, even though we do not believe that they have plans or goals or act with intention.

The words of a story are only surface phenomena of its narrative structure. Equally important is the set of cultural expectations behind a story; they make a story comprehensible to us and allow us to draw inferences from it. When we tell a story we do not mention everything that happened; much is left to implication. For example, if I say that I had breakfast with Mr. Smith at Joe's Cafe, I do not mention every mouthful of food I ate. My listeners naturally assume that we went to a restaurant, that someone took our order, that both of us ordered food, that we ate the food, and so on. We do not speak about such events unless there is a reason to do so. We always understand a story against a background of other expectations that are also organized and stored in narrative form.

Many simple propositional sentences are actually narratives in disguise. Consider the sentence, "Mr. Smith and I discussed the game over breakfast at Joe's Cafe." This sentence not only states a fact; it also tells a story. But it does so only because it implicitly draws on a whole set of cultural expectations—for example, how to have a discussion, how to eat a meal with someone else at a restaurant, what kinds of things one usually eats at breakfast, and so on.

Thus, at its most basic level, narrative structure is a structure of expectations, which are embedded in and connected to larger networks of expectations. These expectations play a dual role. First, they frame our understanding of what is happening. They give meaning to events. We attempt to understand what is happening in terms of expectations we already possess. We recognize patterns of behavior as meaningful in terms of patterns we are already familiar

with. We create a story about what is happening based on stock stories—expected sequences of events—that already lie to hand. Second, the expectations that frame our understanding create the possibility of deviations from what is expected. These deviations call for explanation, and we employ stories to explain them.

Thus behind all narratives lie understandings about what is canonical, expected, and ordinary. These understandings are themselves narratively organized because they are stored in sequences of actions and events: this usually follows that; this is done on Sundays and that on Mondays; this is how you are expected to behave under these conditions; and so on. But these cultural expectations are Janus-faced: storing information in this way simultaneously determines what is deviant, unexpected, and extraordinary in a situation. It creates an agenda for what does not fit our stock of existing narratives and therefore has to be explained. That explanation, in turn, will be phrased in terms of a story that ascribes motivations, intentions, and beliefs to an actor and relies on other stock stories about human behavior.

This is the dual character of narrative thinking: it focuses on, frames, and uses what is expected in human life, and it bestows legitimacy and authority on the expected. At the same time, narrative thinking lets us organize the exceptional and the unusual into a comprehensible form. It allows us to learn by letting us match and reconfigure old expectations in light of new experiences. In this sense, narrative thinking is a heuristic device. It is one of the most basic of cultural heuristics.

Narratives as Norms

Cultural expectations also act as norms. The word *norm* has two meanings—a benchmark of what is ordinary or average, and a standard of what is appropriate. Similarly, the word *normal* can mean what is expected and what is appropriate to a situation. Narrative thought combines these two meanings. Our cultural expectations help us understand what is happening by reference to norms of what is expected in a situation and what is appropriate to the situation. Thus cultural expectations, stored in narrative memory, help frame social reality. As Erving Goffman pointed out, the frame we use to understand events shapes what we believe is happening and what is socially real. Events that seem normal or obvious in one frame become bizarre or inexplicable in another.² When the frame becomes controversial or blurred, it loses its framing character and our sense of what is socially real is disturbed.

Human beings organize their cultural behavior around expectations because this strategy saves effort in thinking and in determining how to act. Much of what we call cultural know-how involves expectations about what kind of situation we are facing and how to proceed in such a situation. Cognitive psychologists call these expectations scripts.3 A standard example of a cultural script is knowing how to order and eat a meal in a restaurant. Situational scripts save us time and energy in figuring out what is going on and what we are expected to do. They are ready-to-hand narrative constructions that we adapt to various social situations. They offer us roles to play and ways to behave. We can think of them as narratives in which we are one of the actors, plots in which we play some of the characters. Scripts do the work of Goffman's cultural frames. They set up expectations about what things mean, and they offer a background against which events and statements can be understood. Because of our restaurant script, when the waiter says to us, "OK, what will it be?" we understand that he is asking about our order and not the nature of the universe.

Much conventional behavior is oriented around such scripts, which is another way of saying that much conventional behavior is organized around coordinated sets of cultural expectations. When people go into a restaurant, they know what is likely to happen and hence they know how to behave appropriately. Moreover, they assume that others will behave in similar or complementary ways. They assume that the waiter will approach them to take their order and not to extract their wisdom teeth. Thus most cultural understanding begins with a postulate of "situational normalcy": unless there are good reasons to the contrary, people tend to behave normally in accordance with the social situation that they believe themselves to be in, and according to the social roles expected of them in that situation.

This rule of situational normalcy underlies Paul Grice's theory of conversational interpretation. His Cooperative Principle is really a baseline of expectations about communication: we assume, without evidence to the contrary, that communications will be brief, truthful, relevant, and perspicuous.⁴ When people deviate from these expectations, they cause us to search for explanations. Because people have departed from the ordinary scripts of conversation, we must make sense of their behavior in some other way.

Social scripts offer background expectations about what is happening, what is ordinary, and what things mean. These expectations literally go without saying, and that is why we do not usually speak about them. That is why our replies to the waiter are brief and perspicuous, to use Grice's terminology. Indeed, it defies cultural expectations for people to attempt to articulate the nature of a script that they are following in detail. Suppose a waiter approaches us and we say: "I see that you are a waiter. You are here to ask me what food I would like. The piece of paper in your hand is a menu. Give it to me and I shall tell you what I would like to order." The waiter would think we were crazy, or obnoxious, or performing some sort of psychological experiment. To talk about frames is to make them lose their character as frames—to make them a possible subject of analysis and contestation, which must be framed by some other set of expectations. Much deconstructive argument—and many artistic effects—involve shifting cultural frames or making us self-conscious about them in order to disturb our sense of normalcy.

What is ordinary about the ordinary is precisely that we don't comment on its ordinariness, don't feel that it needs explanation or explication. Only deviations from the normal are worthy of comment. Cultural know-how, in this sense, is the ability to understand the ordinary and have our expectations confirmed by experience. Conversely, to lack cultural know-how is to fail to recognize the ordinary as ordinary, to lack expectations about what is going on and how to behave. So if a person asked us, "Why did that man walk over to you and give you that piece of paper?" and we thought that she was sincere in asking the question, we would think that she did not understand the cultural norms involved in eating in a restaurant. The remark would be evidence that this person lacked a certain kind of cultural know-how.

Matters are different, however, with acts or events that seem to deviate from the ordinary or the canonical. These things create puzzles that need to be solved or given meaning. Here again narrative structure plays a dual role. Narrative structures offer norms that give meaning to human action, but they also create the possibility of deviations from these norms. People must also be able to make sense of these departures as meaningful human actions. They also use narratives for this purpose.

When we encounter a person who seems to be acting in an unusual or unexpected fashion, and we ask why, we usually get an explanation in terms of a story that ascribes reasons, beliefs, and intentions to the actors involved. That is how one might account for the earlier example in which a customer elaborately described his actions to the waiter: "He told the waiter all these things because he is a social psychologist," or "He said all these things because he is a jerk." Often these explanations are offered in terms of their appropriateness to some other script or set of cultural expectations: for example, how psychologists test people's reactions by doing strange things, how uncouth people tend to tease others, and so on. These actions make sense in terms of these alternative social scripts. Justifications and excuses are familiar forms of narrative explanations. To excuse or justify behavior is to tell a particular kind of story about beliefs, intentions, and actions.

Yet narrative is not only a framework for making behavior meaningful; it is also a framework for understanding the psychology of others and attributing mental states to them. Narrative structures organize our use of psychological concepts like purpose, desire, intention, and belief. When we explain people's behavior through narratives, we simultaneously ascribe purposes, desires, intentions, and beliefs to them: "He ran out of the restaurant because he heard

that his house was on fire"; "He arrived at the meeting thirty minutes late because he wanted to get an advantage in the negotiations." Narratives ascribe mental states to others (or to oneself) to justify or account for deviations from what is culturally canonical or socially expected. Conversely, our ascriptions of belief and purpose make sense because they implicitly rely on background cultural expectations. It makes sense to believe that a person has certain beliefs or desires because of the way she reacts against the background of existing cultural conventions. In this way narratives mediate between beliefs, desires, hopes, intentions, and actions on the one hand, and existing cultural conventions on the other. They understand the former in terms of the latter. In short, narrative structures are a medium through which three facets of human life are understood and explained in terms of one another: (1) cultural conventions, (2) human behavior, and (3) beliefs, plans, goals, and desires.

Narratives as Stock Stories

As forms of cultural software, narrative structures can be passed to others through communication, imitation, or other forms of social learning. Many narrative structures are transmitted through mass media, through artistic expression, and through myths and legends. Much art is based on narrative structures, and our understanding of art is based on absorption and appropriation of these narrative structures. Adults and especially children like to hear stories told over and over again, just as they enjoy hearing a song played repeatedly. Listening to a familiar story can be pleasurable because it fulfills our expectations. Hearing new stories reinforces or alters the existing stock of narrative structures that we use to make sense of what is going on in the world. Art manipulates and plays on our storehouse of stories, retelling them with interesting variations and details. And art can also replenish and expand our stock of stories by exposing us to new narrative structures, new ways of behaving, and new ways of understanding.

Although art is a crucial method of memetic transfer, one of the most important ways that we assimilate scripts and social expectations is through watching other people. The transmission of narrative structures through observation is a good example of how the spread of memes differs from a simple copying of information. Usually people do not transmit social expectations like messages that are coded and uncoded. Rather, watching others acting out social scripts in front of us creates expectations in our own memories. Moreover, because people have different bodies of experience and different sets of prior expectations, they carry away different things from their social encounters. They assimilate behaviors and produce expectations in slightly different ways. As a result, each person in a given culture will have a slightly different set of

social scripts, with slightly different expectations, and therefore each will understand and react to the behavior of others slightly differently. Furthermore, social scripts are not simply routines that we must invariably follow blindly. They are platforms for innovation and improvisation. Precisely because narrative structures give us a sense of the world around us, they enable further development. People play and experiment with narratives and social scripts, producing new expectations that can, in turn, be passed on to others.

People employing similar narrative structures will understand the world in similar ways. Shared social meanings and conventions are not supraindividual entities but result from the interaction of distributions of relatively similar memes. There is no grand restaurant script in the sky, only different but relatively similar restaurant scripts stored in each of us. Nevertheless, these scripts often have an interlocking character. Our expectations about restaurants include expectations about what it is normal to expect from others and what it is normal for others to expect from us. We not only expect that some things are the norm, we also expect that other people also expect that they are the norm. Interlocking expectations can have a stabilizing effect on social conventions and keep them from diverging too widely.

As each of us grows up, we gain a library of social scripts and stock stories. At any time we have an enormous number of stories and parts of stories in our memory. When we want to understand what is happening in society, we try to understand events in terms of an existing story or script. Our recognition of events as an example of an existing storyline creates expectations about how events are to continue. Events may surprise us, and then we try to reinterpret them as following yet another story line. In this way, we try to assimilate what is new in terms of what is old, improvising and playing different stories off against one another to explain deviations from our expectations.

It is important to emphasize the creative aspect of this process. Narrative understanding is not simply a matter of rote; it is also a framework for improvisation and growth. Our library of stories and scripts is constantly increasing. We modify stories and scripts in the light of new experiences; these modifications become part of our memory, used for understanding subsequent events. Suppose, for example, that we go to an Ethiopian restaurant where no silverware is served and people eat with their fingers. Over time we may develop a special set of expectations for Ethiopian restaurants. We may even be surprised if we find a fork on the table at the next Ethiopian restaurant, and we may conclude that the owners are catering to the tourist crowd. (Note how we ascribe motivations to explain deviations from what has now become culturally canonical.) New experiences rewrite our storehouse of narrative expectations, and we improvise on old stories to respond to them. In this fashion our cultural

software is continually rewritten. An increasing variety of narratives adds flexibility to our framing of events and consequently our understanding of them. A person who has "seen it all before" is a person who has many different stories to draw on.

Like other forms of cultural software, new stories are created from older ones through bricolage. Parts of stories or scripts may be combined or grafted onto each other to form new ones. As a result, many of the stories and scripts that we possess bear structural resemblances to one another, even if they are used for widely different purposes. In the same way, we should also expect that many narratives and scripts widely dispersed in the larger culture will be strikingly similar, because they are common descendants of older stories and scripts that have been adapted to new ends.

Narrative understanding is a simultaneous process of organization and matching. To see the present as connected to the past we must already have begun organizing it into narrative form. There is more than one way that one can do this, because every event has many different "hooks" or indices that can connect it to many different stories or scripts. In my memory of a dinner I may recall that the waiter seemed rude, while my friend the oenophile will remember the quality of the wine. I connect this meal to previous stories of rude waiters while she connects it to previous experiences of great wines. Thus it is possible for different people to remember the same events in different ways because each sees its similarities to different kinds of stories and stores it differently in her memory.6

Narrative structures shape our thought because they organize our memory of experience and our methods of memory retrieval.⁷ Narrative structures provide "boxes" into which subsequent events can be categorized, indexed, and stored for later use. Stories and scripts are linked to other stories and scripts through this process. Experiences that do not conform to our existing forms of memory storage are more likely to be lost from memory.8 Indeed, large amounts of our everyday experience are discarded because they do not mesh with our modes of storage. Many aspects of life will be lost to us or remembered in highly limited form if we have a limited stock of stories to serve as an interpretive matrix for categorization and memorization. Just as a pigeon cannot make sense of Hamlet, so a person with only a very small set of stories and scripts will not glean or recall very much information from her experiences.

Not only do we tend to retain memories that conform to our existing narrative structures, we also tend to alter our memories to conform to our canonical expectations about and representations of the social world. Memories that cannot be altered to fit to our expectations may be forgotten or may be deliberately highlighted as exceptions that need explanation. In one famous experiment, college students were asked to tell each other a Native American tale. They either forgot the elements that were unconventional from their own cultural standpoint or transmuted them into something more conventional and expected. This result is hardly surprising. If we don't understand what is going on in an experience, it is more difficult to remember it; it is much easier to understand events in terms of our existing stock of stories. This is consistent with an evolutionary model of memetic development: our minds form an ecology in which certain memes are more likely than others to take root and thrive.

On the other hand, we can, with sufficient effort, gain new sets of cultural expectations. This will change our mental ecology. If we study Native American culture, the previously alien elements of the story gain significance for us and we can remember them when we retell the story. We may even highlight them to our audience as an example of what is distinctive about the culture.

Political and legal rhetoric gains much of its power from these features of narrative framing. People naturally attempt to explain gaps in events for which they have no direct evidence, or events that they do not wholly understand, in terms of familiar stories and scripts. Once they have settled upon a story to frame events, it can exercise great power over their imagination, leading them to make unwarranted inferences and prejudicial judgments.

During the Supreme Court confirmation hearings of Justice Clarence Thomas, Professor Anita Hill accused him of sexually insensitive conduct. Because of conflicting testimony, it was difficult to know whom to believe. Thomas's defenders, however, invoked the plot of a contemporary movie, *Fatal Attraction*, to paint Hill as a spurned lover who was seeking to destroy her former boss's reputation and career. Hill's accusations could then be reinterpreted as those of a calculating, unstable vixen. Thomas himself invoked a stock story about whites who attempt to keep "uppity blacks" in their place, and he accused the Senate Judiciary Committee of staging a "hi-tech lynching." These narrative framings had enormous rhetorical power and may have helped turn the tide in favor of Thomas's eventual confirmation.

Trial lawyers have long understood the power of narrative framing. They attempt to lay out a story of how events occurred during their opening arguments in the hope that the jurors will use the story to frame the evidence they hear. Getting the jury to accept one side's story as the most plausible framework for the events of the trial is often tantamount to winning the case. That is because once a story is accepted, it is used to filter and organize all of the evidence subsequently presented. Like most people, jurors tend to discount or ignore evidence that does not fit their organizing story, and they will alter or simplify information so that it does conform. Evidence that can be made to fit actually tends to reinforce the power of the story, because it seems to confirm it, even though the same piece of evidence could also be consistent with a very

different story. Because narrative framing is so powerful, lawyers faced with the other side's story realize that if they are to win the case they must offer an equally plausible counterstory that also fits most of the evidence. Often the only way to dislodge a narrative is with another narrative that also fits most of the facts but shows them in a very different light.

Personal Narratives

Judgments of human character are also organized around narratives. We form expectations about people's behavior and ascribe attitudes to them that are consistent with these expectations. When we describe what individuals are like, we often do so in terms of stories about the sorts of things they usually do or the kinds of things we normally expect of them. We often judge people and explain them to others through anecdotes that reveal their characteristic behaviors and attitudes. Family members often have stock stories that they tell about other members. These stories not only describe the character of particular family members but also their place in the family, as beloved firstborn son or black sheep. We also use anecdotes to characterize organizations and even entire cultures. Travel guides often contain a wealth of anecdotes that create expectations about how members of a given country are likely to behave.

Just as social scripts are made from fragments of older ones, our expectations about people are created from expectations about other people who seem similar to them. Moreover, in our culture we have a well-developed set of stock characters—the miser, the ladies' man, the clinging mother, the neurotic intellectual, and so on—that we use to frame our understanding of others. These stock characters are templates of expectations that we use as building blocks to form our expectations of particular people that we meet. Although our expectations about people may change greatly as we learn more about them, the initial framing of individuals as fitting a certain stock character may have a significant effect on the future development of our expectations about them. That is because we will tend to behave toward them according to our existing expectations. In this respect, it may well be true that first impressions are lasting ones.

Often our pool of stock characters is tied to ethnic and gendered stereotypes. We have stock stories about how whites and blacks or American tourists and French waiters normally behave (and therefore are expected to behave) with respect to each other. We have literally hundreds of stock stories about how men and women behave, for example, with all sorts of variations—the macho man, the passive wallflower, the stupid hunk, the femme fatale, the henpecked husband, the ditzy blonde. Each of these stock characters can form a template for organizing and giving meaning to our encounters with others.

They are important parts of stereotypical thinking and often have significant ideological effects.

Narrative structures are extremely important in interpersonal interactions, and especially close personal ones. Individuals in the early parts of an intimate relationship often talk incessantly with each other so that they can create a set of expectations about who the other person is and what he or she is like. They create cultural software in each other for mutual understanding. As time goes on they tend to spend less time in this sort of talk because the pictures are starting to become more fully formed and they have less need for new information. Getting to know people is, in a large sense, creating a set of expectations about them and about their behavior, and ascribing attitudes to them on this basis. In the early stages of a relationship, surprises and unexpected behavior may be desirable, because they are opportunities for learning more about the other person and developing new expectations. Later on in a relationship, surprises may have the opposite effect—they may lead to the unsettling feeling that we do not really know the other person any longer.

In times of crisis, people often need to reconfigure their views about each other, and they may start talking a great deal again. One reason people in longterm intimate relationships "never talk anymore" may be that they don't feel that they have to. But when ways of behaving are no longer satisfying-because of a crisis, for example, or because people have grown apart—their old expectations may no longer be adequate. At that point, the fact that people "never talk anymore" becomes a real problem. As a result, people may feel compelled to start intensive discussions about themselves again, in order to reconfigure their expectations about each other and preserve their relationship.

We use narratives not only to describe our personal experiences to others but also to understand them ourselves. Often people feel the need to talk about their experiences to others so that they themselves can comprehend them. The act of talking organizes experience into narrative form so that it can be understood and memorized. The need to talk and describe what has happened may be especially great with regard to emotionally powerful experiences. People must connect these experiences to narrative constructions that they already understand and to features of their lives that they already recognize. This may require considerable narrative work.11

Narrative construction of personal experience is inevitably partial; it selects certain features of experience as meaningful because narratives are organized in terms of what is already understood to be meaningful. What cannot be so organized is usually and eventually forgotten over time. As a result, narrative memories of the same event by different people can vary widely. Each person remembers what is most salient to her, given her existing cultural software and her special preoccupations.

Just as people ascribe purposes and motivations to the behavior of other people so that they can understand it, they do the same with their own behavior. When asked to describe themselves, people often give stories about what they have done in the past and why they did it. They offer anecdotes about their past that symbolize the kind of people they believe they are and the way they usually behave. They describe events that have shaped them and have made them who they are today. Equally important, people construct narratives of their lives. They understand who they are, what is happening to them, and what they should do next by means of narratives. These narratives are stories in which they are the protagonists. Such stories often fit well-established patterns—the Bildungsroman, the adventure story, or the picaresque novel. They portray people's lives as comedy, tragedy, or even farce. 12

Personal narratives organize and give meaning to previous experience. They also provide a form of justification. We justify who we are to ourselves in terms of a story about what we were before and what we have gone through. Perhaps most important, autobiography is a form of prophecy. The presonal narratives—whether comedies or tragedies—can be seen as a kind of script. And scripts are meant to be followed. A personal story has a trajectory, a trajectory that demands to be filled out through future action. A personal story is a set of expectations about the self that demand to be fulfilled in practice. If we see the story of our lives as a tragedy, we may understand what we must do next and what will eventually happen to us in tragic terms. The role of personal narrative in framing the possibilities of our future actions, and thus in limiting or empowering us, cannot be overstated. In extreme cases, we can become the slaves of our personal narratives.

Although autobiographical narratives are deeply personal, they also make use of stock stories and elements available in the surrounding culture. We understand whether we are successes or failures, good or bad persons, in terms of social roles, stock stories, and stock characters. Our own narrative understanding of ourselves is composed out of elements that we get from the larger culture—from movies, television, family anecdotes, social mores, and cultural expectations.

Moreover, our individual narratives are strongly influenced by our cultural heritage. Our ethnic and religious identity forms a template of expectations about how to behave toward others, how to be a man or a woman, how to act toward our children or our parents, and so on. Our understanding of ourselves as Jewish or Italian, Korean or black, already preshapes and constrains the possible stories we build upon and the kinds of futures we feel that we can have. These effects on personal narrative are another example of the inextricable relationship between the personal and political, the individual and the cultural.

People use and develop scripts for their interactions with others. Just as people learn how to order meals in a restaurant by watching others, so, too, they learn how to get along with others, handle and avoid conflicts, love and be loved by watching their parents and others close to them. They develop narrative expectations for how to be a friend, a lover, and a parent. They modify and rework the scripts they have learned in previous social settings to form scripts for dealing with people in new situations. Our expectations about social relations with others are produced through bricolage from previous relationships.

People learn how to be parents, for example, from watching their own parents; this learning helps them create roles that they naturally slip into in their dealings with their own children. People absorb lessons about how to deal with and love others from their parents and others close to them; they apply these lessons to their subsequent relationships with others. These scripts may be particularly awkward examples of bricolage—old tools badly adapted to fit new social situations. Yet people cling to these scripts because they do not know how else to perform these roles.

We often employ the narratives of others consciously or unconsciously as models for our lives. We may absorb the stories of our parents into our own personal narratives, for example, using them as the raw materials to develop our own personal stories. As a result, we may feel unconsciously compelled to play out parts of these stories in our own lives. A man whose father failed at business may absorb this story into his own personal narrative and reenact it as part of his own life story. Although we are hardly doomed to repeat the narratives of our parents in exact detail, parts of their stories may still be important elements in what we eventually do construct. And, as is so often the case with bricolage, even our modifications and innovations may bear the characteristics of what was built upon.

The narrative nature of human self-understanding and human social interaction explains why psychological therapy has historically turned to narratives as a means of treating patients. Much psychological therapy involves recounting, interpreting, and reconstructing the patient's story and critically examining the scripts she employs. Through recounting and revising stories about herself, the patient begins to recognize how she became the person she is today. She learns to identify the sources of the scripts and expectations that underlie her reactions to people and events. She tries to see how living according to these scripts and expectations is keeping her from a happy, healthy, life. Together the patient and therapist try to modify her scripts and expectations by substituting new narratives for old ones.

In short, successful therapy teaches the patient to develop new scripts and write new cultural software through repeated narrative construction by patient

and therapist. The patient learns to reorganize the past in new ways, to see previous events in a different light, and to form new expectations. The hope is that these new expectations will lead to new and more positive behavior.

From a purely physical standpoint, it has often seemed puzzling how the mere recitation of stories could effect any improvement in a person's mental condition. But if people's behavior is shaped by the narrative constructions that they use to understand themselves and interact with others, if much social thinking relies heavily on scripts and expectations, the idea of a "talking cure" is not at all far-fetched. Supplementing or replacing old narratives and scripts with new ones might change people's behavior for the better. The problem is that the means of doing this must necessarily be as much an art as a science. And it is by no means clear that there is only one route to the creation of new and healthier scripts.

There is an important analogy between the personal narratives and scripts that hinder our development and make us unhappy on the one hand and the ideological effects of cultural software on the other. Just as people's cultural software contributes to social injustice, it can also contribute to their personal unhappiness. Indeed, some conceptions of justice do not draw a sharp distinction between these two concerns. Under those conceptions, the idea of justice also applies to the self, so that one can speak of people being unjust to themselves. This injustice is not a matter of bad behavior toward others (although this can be involved). It is rather an injustice to the possibilities of what we could be.

The narratives of our lives and the social scripts that we employ in interacting with others can lead us repeatedly to act in self-destructive ways or in ways that prevent happy and fulfilling relationships with others. We may seek out lovers who abuse or manipulate us, for example, because we are replaying scripts about how to love and be loved that we began organizing in early childhood. We may deal with conflict in ways that we worked out for situations that happened long ago, strategies that are inadequate to the situation that now faces us. We may generate unhappiness and conflicts with our children because we are following scripts of how to be a parent that we assimilated from our own parents.

These personal scripts and narratives and their unfortunate consequences for our lives are like the ideological effects of cultural software, except that they act at a very personal level, and our concern with them is not that they produce social injustice but that they hinder personal growth and personal happiness. Just as we must take an ambivalent attitude to our cultural software because it has the capacity to produce injustice, we must take an ambivalent attitude toward cultural software because it has the capacity to produce personal unhappiness. Personal scripts and roles that we have assimilated may be partially adequate to deal with many of the problems and situations we face in our lives. That is probably why we developed and adopted them in the first place. But in new contexts and new situations, their inadequacies become increasingly apparent, and they begin to hamper our lives. The goal of successful therapy is to build newer, more adequate expectations out of older, less adequate ones.

Group Narratives

Just as individuals have stories that they use to understand themselves and the world around them, so do entire cultures and countries. Each society has stock stories drawn from its past that are told over and over, and, in this retelling, take on a mythic status. These stories symbolize what is most important to the society, its values, its sense of itself, and its relationship to the outside world. The stock stories of a society are abstracted and condensed through frequent retelling, and eventually can even be encoded as single icons, which can be either persons or events. Examples in American history are "Lincoln," "Pearl Harbor," or "the Alamo." Each of these icons condenses and thus represents a story with a rich set of historical associations, often contradictory and contested. Such narrative icons invoke not only a particular order of events but also a tradition of interpretations that grows up around these events.

These stock stories and icons form part of what is variously called collective memory or social memory. It is an excellent example of widespread memetic transfer and assimilation. Social memory is an example of an endemic cultural virus. Narrative memory is spread from generation to generation through communication and becomes part of the shared cultural software of a culture or society. Members of a society or culture repeatedly tell each other stories about important events in their history. These events often have deep emotional resonances; they are still able to evoke anger, solidarity, pain or pride long after the events have passed. Examples are wars and revolutions, depressions, riots, strikes, famous trials, and genocides. These important events become benchmarks for comparison with later events, an index through which to understand what is happening to the culture. As the memories are passed on through the generations, they are stylized, pared down, and altered, much like any other story. These social memories are thus at the disposal of storytellers, the mass media, and even the state, to rework and reorient.

People share social memories because they are members of a common meme pool. As a result, entry of new individuals into the meme pool can alter social memory; even though widely shared memories can be assimilated into the new members, they may also bring with them new stories from different cultures.

Social memory is distinct from other endemic cultural software in an im-

portant way. The memory of important events is not simply shared by members of a culture. It also unites them or divides them, gives them something in common or produces a bone of contention. Some especially divisive events, like the Vietnam War, the Dreyfus trial in France, and (most likely) the O. J. Simpson trial in the United States, retain their ability long afterward to invoke conflicting meanings and reinstitute old social and ideological divisions.¹⁶

Sometimes the divisions created by a momentous event will be resolved or mediated by the creation of a stock story that is roughly satisfying to most of the contending groups. American culture has produced such a narrative of the Civil War: the North fought for freedom, but the South is acknowledged to have fought valiantly and bravely, and the conflict is seen as deeply tragic. The noble figure of Robert E. Lee-who gave up the chance to command the northern armies in order to defend his homeland of Virginia-plays an important role in symbolizing southern heroism against enormous odds, as well as the deep, familial connections between the two sides even in the midst of a vicious war. This narrative allows both sides to accept and even celebrate the outcome of the Civil War-through the staging of mock battles and the collection of memorabilia, for example. Of course, this narrative is most appealing to American whites, since it sees the conflict as one between honorable brothers who are eventually reconciled. It operates only by downplaying the evils of slavery and the subjugation of an entire people. Thus it is not a completely effective narrative resolution. And it breaks down precisely in moments of heightened racial awareness.

Just as a person draws on a stock of stories to frame and understand what is currently happening, so members of a culture draw upon its stock of stories or myths to frame and understand what is happening to them. These myths organize experience and the culture's reaction to that experience. The value of myth is that it helps us understand what is new in terms of what is already understood. As we saw previously, stories save time and energy in figuring out what is going on, what is socially real, what it is appropriate to do, and what is likely to happen in the future. The myths of a culture reorganize the world to take on the appearance of a story that is already well understood. Thus an attack on American interests is readily envisioned as another Alamo or Pearl Harbor. If the new experience fits the old story well enough, the myth will be strengthened and confirmed. But experience is often recalcitrant. Although we try to understand social events in terms of stories that lie to hand, the world will not always so easily conform. Viewing every conflict as the moral equivalent of Pearl Harbor and every war as World War II is a recipe for disaster. So a culture's myths are gradually adjusted and reconstituted in order to take account of the changing world it faces. The stories of the past are given new

glosses and new meanings. In this way old icons like "the Winning of the West" or "the Civil War" take on multiple and conflicting associations. Widely shared stories serve as a springboard for future improvisation. Like other memes, myths are transformed as they are communicated to others in new circumstances. Thus the myths of a society are not simply handed down; they are subtly remade in each generation for its own purposes, even as that generation is guided and shaped by what it works with.¹⁷

Occasionally new stories and icons replace older ones. "Vietnam" has become a permanent fixture in American consciousness, with a highly conflicted and ambivalent set of meanings. Yet as with all cultural bricolage, the new is always built on the old. A culture's new stock stories are constructed out of variations on older narratives, which in turn were based on still older frameworks originally used to understand the problems of the distant past. Thus it is hardly surprising to find that the stock stories of different cultures have many elements in common, even if these elements have now come to be used for widely different purposes and have quite different meanings.

Because a culture's stock stories are a shared way of understanding the shared past, present, and future of its members, media of mass communication take on a special importance in the creation and development of social myths. They are a crucial determinant of the ecology that narrative memes face. The stories and genres produced by mass media are evidence of a culture's most pressing and abiding concerns. Yet the products of mass media cannot be taken as the endpoint of analysis. They are not simply and faithfully absorbed by members of a culture. Like all memes, narratives take root depending on the existing cultural software of the people who absorb them. Viewers and listeners appropriate and reinterpret the narratives produced by mass media. They are the raw materials for the work of cultural memory performed by individuals. And the sum of these individual appropriations and reinterpretations of stories, in turn, shapes the way that the mass media will present these stock stories in the future.

Richard Slotkin has argued that a central myth of American consciousness is the story of the frontier: a story of repeated separation, regression, and regeneration through violence. According to this story, Americans leave civilization—whether Europe or (later on) the cities of the Eastern Seaboard—and spread out into the wilderness, where they find that they must live more primitively and simply. There they encounter savages who threaten their survival and whom they must conquer. American progress comes from leaving the old world behind, temporarily regressing to a more primitive or natural state, taming the wilderness, and defeating the savages. To succeed, Americans must both separate themselves from the older authority structures and class privilege of

civilization and defeat the savage of the western wilderness. The familiar icon of "Cowboys versus Indians" is only one of the many variants of this central myth.

In this account, the American hero is one who confronts the wilderness and makes it his own. The American hero is one, like Natty Bumppo or Daniel Boone, who "knows Indians" and absorbs this knowledge to transform himself and civilized society. For these American heroes, the war against savages is mirrored by a struggle in their own souls in which they conquer and discipline the savage side of their own natures.¹⁹

The story of the frontier has several corollary stories, including the myth of bonanza: even though the frontier is dangerous, it is often a place of unbridled opportunity in which great wealth can come from comparatively little effort. Freedom comes from conquering the frontier and making use of its manifold opportunities. Repeatedly, events in American history—from the California Gold Rush to the Roaring Twenties to the deficit-financed boom of the 1980s and the hi-tech expansion of the present—have been conceptualized in terms of the boundless wealth of the frontier. In the myth of the bonanza, all things are possible and everyone can grow rich, at least until the gold reserves dry up, the stock market crashes, or the savings and loans go bankrupt.

The myth of the frontier has been used repeatedly by Americans to understand who they are as a nation and how they should behave in the many crises that have faced them. The story of the "savage war" has been used not only to justify the expropriation and extermination of the Indians but also to understand many other crises and problems, both domestic and international.²¹ Labor conflicts, race relations, the Cold War, and the war in Vietnam have all been readily envisioned as new versions of the mythical "war against the savages," in which annihilation of the enemy is the only acceptable way to preserve American civilization. The conception of the savage war projects the difficulties of American life outward onto some hypothesized Other—whether it be Native Americans, labor unions, immigrants, international communism, or the contemporary black urban male—who is seen as a dangerous element that must be defeated or controlled, and "who becomes the only obstacle to the creation of a perfect republic."²²

One might contrast this stock story of American life with the stock story of Judaism. Because Jewish culture is much older, it has a wider range of narratives to draw from. Nevertheless, the most enduring stock story of the Jews is a cyclical myth of dispersion, persecution, and redemption. It is, roughly speaking, the story of the Exodus. In this story, God reveals himself to the Jews and promises them his everlasting protection if they will obey him and spread his wisdom to other countries. The Jews are then dispersed from their

homeland (usually because of their previous misbehavior) into foreign countries. They are, as Moses says when he flees into Midian, "stranger[s] in a strange land." Like Joseph, they benefit their adopted lands and rise to prominence. Eventually, there arises "a new King... who knew not Joseph"—who does not understand the value that the Jews bring to civilization. The new political forces persecute the Jews: "He set over Israel taskmasters to afflict them with burdens. And he made them serve with rigor." God hears the cry of the Jews and remembers his promises to them. He delivers them from bondage "with a strong hand and an outstretched arm, with signs and wonders." He then reveals himself to them, gives them his Law, and delivers them to their Promised Land. This story is implicitly cyclical. The Jews begin and end in their homeland. Once in the Promised Land, they can then be dispersed again for their misbehavior, and the cycle begins anew.²³

This story has enormous pull over Jewish culture and thought. Many Jewish holidays are organized around the myth or various stages within it. The most obvious is Passover, when Jews are directed to tell their children the story of the Exodus. The Passover Haggadah even instructs them that they are to regard themselves as if they personally had been delivered from Egypt. Two other major festivals, Succot and Shavuot, celebrate different aspects of the redemption story—the dwelling in the desert and the reception of the Law. The story of Mordecai and Haman in the Book of Esther is yet another version of the myth, as is the story of the Maccabees celebrated at Chanukah. The fast day of Tisha B'Av commemorates the destruction of the First and Second Temples and the beginning of two different diasporas. The most recently added Jewish holiday is for the commemoration of the Holocaust. In one sense the Holocaust fits well into the governing myth of Judaism, because it is a catastrophic example of oppression, which is followed by the founding of the State of Israel. In another sense it severely strains the myth, because the sheer scope of the persecution involved in the Holocaust seems to dwarf any previous misbehavior by the Jewish people and any promise of eventual redemption, even the creation of a Jewish homeland. The Holocaust is an example of a historical event that has tested and reshaped the fundamental myth of a culture.

Each of the myths just offered is derived from previous sources. For example, the American myth of the frontier probably borrowed from the story of the Jews in the wilderness who gain the land promised to them by God. Note, however, that these two stories take on very different trajectories and that the wilderness serves very different purposes in each of them. If anything, the myth of the American frontier seems to borrow most heavily from the story of the Jewish wars to conquer Canaan and the divine injunction to destroy the Amalekites.

The great danger of myths, like personal scripts, is that they not only frame

our understanding, they also invite us to play them out in our lives. Narratives are not only tools of understanding but also tools of action. As heuristics, they save time in understanding a situation because they provide ready-made social meanings to events and ready-made roles to play in them. Narratives are scripts, and scripts are made to be followed. Yet cultures are by no means destined to play out their constitutive myths over and over again. Any culture that has existed for long periods of time has many different stories and myths to choose from. Moreover, social myths are not simply a script to be followed blindly; they provide opportunities for reinterpretation and a platform for innovation. Each generation has reread and rewritten the dominant stories of its past to serve the needs of the present. Social myths lend themselves to this reinterpretation precisely because they have such deep resonances and such a wealth of associations. People can and do draw on countermyths; they can reinterpret, reread, and rewrite existing stock stories to meet present-day concerns and crises. For example, Americans have not only the myth of the frontier but also the story of Thanksgiving, in which the Indians aid the Pilgrims and the two share the bounty of the land and jointly give thanks to divine providence. This version of the Thanksgiving myth, one might think, would be much more helpful to a country that is trying to live down the sorry history of previous persecutions and trying to accommodate people of many diverse cultures, than is the myth of the savage war.

These examples show how narratives function as cultural heuristics. They help constitute much of our social understanding and make possible much of our social existence. Yet at the same time it is clear that they can have serious ideological effects. They can produce stereotypical thinking or lead us to reenact them in wholly inappropriate situations. Hence our attitude toward them, as toward all cultural software, must be ambivalent.

Moreover, as heuristics, narratives are necessarily partial. They can describe and store in memory only certain features or aspects of a situation. The world is too complicated and multifarious to be captured in a single narrative account or even in a series of narrative descriptions. This inadequacy of narratives is the flip side of their advantage to us. Narratives are useful memory structures precisely because they select and organize our experience—they categorize and store events into scripts or indices that we can use for later comprehension and comparison. Narratives are useful tools of understanding because they create social expectations that frame our understanding of what is and should be happening; without such expectations, we literally would not know what to expect. Because narrative structures work in these ways, they necessarily lead our understanding in some directions rather than others. They categorize future experience in terms of preexisting indices and expectations. This produces the familiar trade-off of any heuristic—although these expectations may be good

enough for some purposes, they may seriously hinder our understanding and promote injustice in others.

Narratives and Justice

Narratives are intimately connected to questions of justice. We use narratives to describe human plans, goals, and intentions, which are often necessary to judgments of what is just and unjust. Moreover, our explanations of what is just or unjust in a situation often depend on a narrative account of how that situation came to be. Narratives connect—and sometimes fail to connect—the misfortunes and inequalities of the present to the events of the past. The fact that my wallet is empty looks very different depending on whether I spent all my money or it was taken from me; my injuries have a different meaning depending on whether I have tripped over something or have been kicked.²⁴

To be sure, theories of justice need not be based on the events of the past. Under some theories of justice, we can judge something as just or unjust simply on the basis of the existing distribution of resources. But the explanation of why that distribution is unjust inevitably will rest on some narrative account that describes either the meaning or the consequences of the distribution. For example, one might hold that a distribution of resources is unjust because it conveys a certain meaning about the worth or dignity of the individual or because it is likely to prevent equal chances for human happiness in the future or is likely to have oppressive or wrongful consequences. In each case, we still offer a narrative account of the evils of the present state of affairs. Our sense of justice inevitably has a narrative character, whether our concerns are corrective or distributive, whether our theory is deontological or consequentialist, and whether our vision of justice is forward or backward looking. Indeed, the very fact that our notion of justice looks in any direction at all means that some narrative underpins our accounts of what makes a situation just or unjust.

Because of the connections between narrative and justice, social memory is an essential framework for judging questions of justice and injustice, and control of social memory is an important ideological tool. The history that has produced present holdings and present injuries is usually important in assessing whether people have been treated justly or unjustly. If people forget the past or if it is disguised from public view, the world seems a blank slate or a level playing field: past injustices are forgotten and present debts are wiped out. A person's or a group's social situation and respective life chances are more easily seen as no one's fault in particular and as primarily the responsibility of each person or group. If narratives of previous injustices are forgotten, distorted, or replaced by false narratives of reconciliation and recompense, serious injustices may go unacknowledged and unremedied. Loss or distortion of social memory

can bury past injustices and make present distributions of power, wealth, and other social goods seem unfreighted with previous wrongs. It is no accident that the words *amnesty* and *amnesia* come from the same Greek word meaning "to forget."²⁵

Yet the memory of previous injustices is not always an unqualified good. If past injustices become deeply and pathologically lodged in our identity, they can adversely affect how we deal with others. Fixation on the memory of past wrongs may be necessary to preserve the memory of unremedied injustices or as an important lesson for future generations, but it can also hinder our personal growth. Both forgetting and remembering can be pathological.

Narratives and the Canonical

Narratives can also have ideological effects because they shape our sense of what is canonical and normal, and hence our understanding of what is different and deviant. Our sense of the canonical includes not only expectations about behavior in restaurants but also expectations about how blacks and whites or men and women are likely to behave. Race and class relations are organized around stock stories about members of different groups. Debates about welfare policy and immigration are often based on anecdotes about the behavior of welfare recipients and immigrants.

Expectations about social groups have both a descriptive and normative character; they can describe not only how things usually are but how they are expected to be. Situations that conform to canonical expectations require no special explanation or justification. But deviations from the canonical—like a black man walking in an all-white suburb or two men kissing in a public park—call attention to themselves as unusual and special and demand either explanation or justification.

Moreover, people often understand majority or superordinate groups as canonical in both descriptive and normative senses. Racial minorities, women, homosexuals, and the physically and mentally disabled are generally viewed as being "different." This "difference" arises from an implicit comparison to canonical norms of whiteness, maleness, heterosexuality, and lack of disability.²⁶

Understanding people as "different" often means identifying their differences from the canonical as the cause of any disruption of settled expectations. What is canonical forms the baseline of expectations; because it is regarded as normal, it is not in need of special justification. Thus a person in a wheelchair cannot climb stairs to enter a building, while "normal" people can. The cause for this situation is located in the disability of the person in the wheelchair, not in canonical social expectations and settled social institutions. Accommodating disabled people may require adding ramps, widening doors, and provid-

ing elevators. Such accommodations may be understood as special treatment even when they are done in the name of equality. The notion that equality entails "accommodation" or "special treatment" preserves a sense of what is canonical even as it authorizes deviations from it. It locates difference—deviation from the canonical—in the minority or subordinate group rather than in the social relation between groups.²⁷

Our narrative expectations about what is canonical and ordinary are far more than isolated sets of beliefs about men and women or about whites and blacks. They are forms of cultural know-how in the truest sense, for they offer us scripts about how to possess a particular identity in a culture and expectations about what to expect from people with different identities. Cultural know-how is in large part the ability to understand what is canonical and to have our expectations confirmed in social practice. This is not necessarily an unmitigated good. Some of these expectations are essential mechanisms in the preservation and reproduction of unjust hierarchies. Considerable cultural knowlege is necessary to behave according to the dictates of a status regime, including an unjust one. One of our most finely tuned social skills is the skill of understanding our status relative to others and the consequences of that comparative regard. This skill is used over and over again in our everyday interactions.

Children are taught how to operate within status regimes through social learning and imitation. They absorb new cultural software that makes them socially competent to be racist or homophobic. People who lack the requisite cultural software may commit social faux pas because they do not understand the social meaning of situations; they do not discriminate in the same ways that others will.

In the previous chapter I compared racism to a virus or disease transmitted through social learning. But such a "virus" is also a form of cultural know-how. Hence one must also understand racism as a social skill, often deeply ingrained in the cultural software of individuals. It is not a lack of cultural knowledge but a particular manifestation of it. We often say that racism is born of ignorance, but in another sense this is hardly so. The truly ignorant person is the one who does not understand the system of racial caste and therefore does not know how to behave within it. To participate in a regime of racial status requires delicate and complicated social understandings. It demands a considerable degree of fluency in the language of injustice.

Narratives That Make Themselves True

Narratives do more than simply distort or limit understanding. They also have the ability to "make themselves true" through their use. Because social meaning is part of the social world and is constructed in part by narrative understandings, narrative organization is folded into the social world, becoming part of its fabric and shaping its future evolution.

First, shared narrative structures help create intersubjective social meanings. People use narratives to understand the nature of a social situation; when many people share the same set of expectations about what is happening—for example, when they employ similar conventions and similar assessments of social meaning—their interlocking expectations establish what is the case socially. If people in a society share the view that people with darker skin have lower status, for example, then people with darker skin do have lower status in that society, whether or not this is just.

Second, narrative structures reconfigure and add to the meaning of past events. Narratives cannot change the past, but they can change how people remember the past and what the past means to them. People remember the past in terms of a set of narratives. These narratives bestow meanings to past events that the participants in those events may not have shared. The American Revolution is a good example. The Revolution, like any other part of the past, does not arrive in a premade narrative package with premade social meanings. The meaning of the Revolution must be constructed out of the memories and stories of the persons who participated in it, those who witnessed it contemporaneously, and those who come afterward. These narratives do not exhaust what happened during the Revolution. They are partial in both senses of the word—they are both incomplete and biased in their organization and characterizations. Yet they add something to the social meaning of the American Revolution. The narratives produced about the Revolution become part of the understanding of the people who lived through those times as well of those who came later. These understandings are then passed on to others, who in turn add their own interpretations and stories about these events.

In this way a tradition of understandings and narrative expectations about the American Revolution grows up. This tradition is always in flux, shedding old meanings and gaining new ones. Its associated memes develop and mutate as the tradition is passed on from person to person and from generation to generation. The resulting narrative structures are folded into the tradition and become part of the social meaning of the American Revolution. They need not form a homogenous whole and may be in conflict with each other. As a result, the social meaning of the American Revolution may be contested and contradictory. Nevertheless, the cumulative social meanings of the Revolution are quite real, even if they embellish or mischaracterize events. The existence of a tradition of representations and the palpability of its effects are distinct from the accuracy of the representations themselves.

Third, narrative structures can make themselves true in practice because they subtly direct the actions of people who frame experience according to these narratives. People tend to characterize situations that they face according to the scripts or narrative accounts that they possess. These scripts not only describe the nature of what is going on but also offer the nature of a proper response. Depending on the narrative structures that we possess, the same behavior can be interpreted as a mere social slight, a misunderstanding, an aggressive action, or a vital threat to national security. This interpretation can shape our response; the response, in turn, can induce behavior from others that confirms our worst fears.

Antagonistic cultural groups and nations often understand each other's actions according to a previously prepared script of expectations that both reflects and reproduces distrust. Their internalized story lines about the social meaning and the likely course of their relationship may exacerbate tensions between them. Their expectations may lead to mutually self-destructive behavior unless the parties learn to interpret each other's actions and intentions in different ways.

Stock stories are one of the most powerful sources of social and personal prejudice, not merely because they frame the nature of events but because they are mechanisms of self-fulfilling prophecy. People with particularly powerful stock stories have expectations that tend to dominate their interactions with others. They tend to make the world fit these stories by understanding events as confirming examples of their powerful and well-worn narratives. Seeing the world and the behavior of others in this way tends to place others in the roles designed for them. In this way scripts can sometimes "make themselves true." If we see the actions of others continuously in terms of expected slights to our ego, for example, we may well behave in ways that bring about what we fear most. People whose personal scripts end with rejection or abandonment may orchestrate events so that rejection or abandonment becomes likely. If we have stock stories in which a certain group of people is worthless and undeserving, they are likely to treat us with disrespect and hostility, confirming our bad views of them. Nations whose social memory is organized around certain forms of conflict and defeat will often find that they create the very sort of enemies that they expect and deserve.

Of course, no script, no matter how powerful in our imaginations, can completely reorient the behavior of others. But others are also attempting to understand our behavior according to a set of narratively coded expectations that they already possess. Hence our responses to them can often redirect the ways they are likely to respond to us. The people whom we treat as threatening may not become threatening simply because we expect them to; but our aggressiveness toward those people may cause them to act aggressively toward us in return, which then confirms our estimation of their dangerous tendencies.

I noted earlier that a familiar set of American cultural expectations under-

stands opposition and adversity in terms of a "savage war" in which Americans must dominate and conquer their enemy. Such a script followed to its logical conclusion often tends to provoke violent responses from the group that is assigned the role of the "Indians," thus providing the country with the conflict it expects and deserves. Yet this behavior may have disastrous consequences, as in the case of the war in Vietnam: Although the United States responded to the crisis in terms of its traditional script, the Vietcong were not Indians, the American soldiers were not Cowboys, and the result was not the winning of the Old West.

Ironically, narrative framing can make itself true because narrative structure is irrelevant to truth. What is true and what might be true are both expressed, understood, and memorized in narrative form. In addition, narrative structures are both a set of frames for experience and a set of directions for action. Narrative structures do double duty in social memory and social convention. The distance from explanatory story to cultural script is not very far, because both are constructed from the same cognitive materials.

Because both history and conventions for social behavior are stored in narrative form, there is an inevitable tendency for the two to nourish each other and be confused with each other. The importance of history to the human mind is precisely its tendency to make us want us to reenact it, to follow its lead, to see the path ahead of us in terms of the path that was lately trod. Professional historians deliberately resist this impulse—for they are interested precisely in discovering and showing the strangeness of other times and lands. But this is an acquired tendency that not even the professional historian can fully adhere to. History inspires; it inspires us to reenact it, to see its relevance to our own time. It presses its events and expectations on us like a dancer whose bodily movements entice us to imitate them, like a musician whose playing energizes us to beat time to its rhythms. History's narrative construction draws us subtly and inevitably into a web of imitation and mimicry, a conflation of history and script, memory and expectation. Santayana had it precisely backward, for it is those who learn from history, who absorb the narrative structures of the past, who are most drawn to and destined to repeat them.

Even the student of history who studies it to avoid its mistakes, as Santayana suggested, is drawn into this web. For when we learn from history, we still engage in mimicry. We imagine ourselves at the scene of the battle so that we can remake the fateful decision. We view ourselves in the situation where the mistake was made, and not in a wholly new one. We see a law that binds together the past situation and our present one as of the same general sort. We postulate a cause and effect that occurred in one setting and that will occur again if we do not choose otherwise. In this way we are still repeating history,

still framing our expectations about what will happen in terms of the narratives of the past. We are simply trying to tell the story differently from a certain point on. And the scripts of the past still have a hold on us. For they suggest that it was this mistake that we should learn from, and not another—one that might have occurred to us if we did not compare our present situation to that particular one but to yet another not in view. Perhaps we think we will not make the same mistake as Napoleon at Waterloo, but perhaps we have already made a mistake in thinking that Waterloo is the appropriate analogy.

Yet even as our expectations are played out, they are in the process of changing. We do learn from experience; that is the flip side of our ability to memorize and repeat. So our myths and stories mutate partly in response to recalcitrant experience, an experience always mediated by our narratives and frameworks as well as those of the others we interact with. Narrative shapes the way history will proceed, but the procession of history is absorbed into memory and reconfigures our expectations and reconstitutes our governing myths. Americans may always think of themselves as Cowboys battling Indians, but the Vietnam War was not the Wild West. Nor was it even Custer's Last Stand. The war in Vietnam has reconfigured American memory and American myth in important ways; it has become its own cultural icon: Vietnam. Just as Americans remind themselves to remember Pearl Harbor, they now assure themselves that there will be no more Vietnams. Thus when Americans fought the Persian Gulf war against Iraq, they developed the military doctrine of "overwhelming force," so that they would not suffer yet another defeat at the hands of the Vietcong. America has surely learned lessons from Vietnam; what is unclear is whether it has learned the right lessons.

The more pervasive and powerful a form of cultural software in understanding the world, the more pervasive and powerful its potential ideological effects. Narrative thought is an excellent example of this phenomenon. Because narratives are so central to our thinking, they create particularly compelling ideological mechanisms. Moreover, narratives produce ideological effects not only because they present a partial or misleading picture of the social world but because they are ways of intervening in the social world and of influencing the responses of others. Narrative structures do not simply reflect the world badly, they shape the world to their own distorted lens. They are not only illusion but prophecy. It is as if one could make one's face become ugly by looking at it repeatedly through a funhouse mirror. Indeed, the optical metaphors of distortion are entirely inadequate to describe the variety of ideological effects that narrative thought can have on the social world.

Because language and symbol are the most pervasive forms of cultural transmission, they offer a rich trove of cultural software and ideological effects. This chapter examines the cultural software that is used to form social meanings transmitted through language in the form of conceptual oppositions and networks of conceptual associations.

Much political and social reasoning draws on conceptual oppositions. A simple and powerful example concerns American attitudes about race. American culture understands whiteness and blackness as opposites, even though there are several races in the United States, and the boundaries between those groups are hardly distinct. More important, the opposition between whiteness and blackness is understood in terms of a network of evaluative conceptual oppositions: law abidingness as opposed to criminality, morality as opposed to immorality, higher intelligence as opposed to lower intelligence, knowledge as opposed to ignorance, industry as opposed to laziness, and so on.¹ This network of associations is an important ideological mechanism in producing an image of black Americans as an inferior Other onto whom all manner of unsavory characteristics are projected. Moreover, this network of cultural associations helps sustain unjust stereotypes about whites and blacks, wrongly presents them as separate, homogeneous, and unified groups, and helps perpetuate misleading assumptions about American society. For example, many American whites associate drug use, welfare dependency, and crime with blackness, and politicians regularly play on these associations, sometimes subtly and sometimes not so subtly. Nevertheless, a majority of drug users, welfare recipients, and criminals are white.2 Thus the network of cultural associations projects criminality and immorality onto blacks and away from whites. This has a dual ideological effect,

identifying blacks as the cause of America's moral problems, and making white criminality and immorality relatively invisible.³

Semiotic Systems as Cultural Software

To explain how networks of association operate, I will use a well-known theory of cultural meaning—structuralism. Structuralism argues that cultural meaning is produced when subjects understand the social world through conceptual oppositions. These oppositions may be between things (sun and moon), directions (left and right), abstract concepts (reason and passion), degrees or qualities (higher and lower), classifications (male and female), and groups of persons organized by gender, race, ethnicity, or social class (men and women, blacks and whites).

People make sense of the cultural world not through isolated conceptual oppositions but through networks of linked conceptual oppositions. When people understand two things or concepts as opposed, they relate this opposition to other oppositions they are already familiar with. The difference between A and B is understood in terms of the difference between C and D. A's association with C is mirrored by the association of B with D. New conceptual oppositions are fashioned by analogy to, and understood through association with, previous conceptual oppositions. The process of concatenation and collation of conceptual oppositions goes on indefinitely. In this way, a huge network of associations develops that produces rich sources of cultural meaning.

This approach to cultural meaning is the basic insight of structuralist theory; it also underlies the semiotics (or semiology) of later thinkers like Roland Barthes.⁴ For this reason, I shall speak of structuralist and semiotic analysis interchangeably. By each I refer to the basic strategy of studying cultural meaning in terms of networks of differences and conceptual oppositions.

Nevertheless, structuralism is also often associated with two more controversial claims about human cognition. The first is belief in innate ideas (associated with Noam Chomsky); the second is belief in universal structures of the human unconscious (associated with Claude Lévi-Strauss). The idea that cultural meaning is constructed through networks of conceptual oppositions, however, requires neither of these assumptions, and the theory of cultural software is committed to neither. To the contrary, I shall argue that the structuralist or semiotic method does not identify stable or permanent conceptual oppositions, though structuralist theory often presents them as such. Rather, what this analysis identifies are the traces or marks of human cultural software as it evolves and combines in different cultures.

When Lévi-Strauss offered his views about innate human ideas, he attempted to solve the problem that faces all accounts of cultural understanding:

to explain why and how shared meanings are shared. His solution, like Chomsky's, was Kantian-style: every person possesses the same unconscious structures; hence each has the same tendencies to understand the social world in terms of networks of conceptual oppositions. Nevertheless, Lévi-Strauss's account does not really explain why particular substantive ideas are opposed to each other, why we find different combinations of ideas opposed or connected to each other in different cultures, or why people in the same culture might have different sets of cultural associations. Nor can his account explain how networks of cultural associations might change over time. At best he offers a formal and ahistorical account of cultural understanding.

Moreover, his assumptions are largely jettisoned in later versions of semiotics and semiology: semiotic studies of advertisements and the fashion system, for example, do not presuppose universal unconscious structures. They are concerned with codes and meanings that change historically, often quite rapidly.⁵ But when semioticians ignore Lévi-Strauss's assumptions about the human unconscious, they simply beg the question of how these cultural codes are created and shared.

A distinct advantage of the theory of cultural software is that it can offer answers to these questions. The human mind readily absorbs and memorizes conceptual oppositions and networks of conceptual association, just as it does in the case of narrative structures. This tendency may be innate, or it may be a side effect of other very basic forms of cultural software. Whatever the cause, networks of association can be and are widely dispersed through communication and social learning. For example, the cultural associations of whiteness and blackness that I mentioned above are stamped into innumerable cultural symbols and messages that Americans absorb from their very earliest years. They are woven into countless sets of expectations and social meanings that we depend on in our daily encounters with others. That is one reason why racist forms of thought are so powerful, pervasive, and difficult to eradicate. Nevertheless, these associations have not always existed; they are the products of memetic evolution. And one day, we may hope, they will again cease to exist in the minds of human beings.

The theory of cultural software also accounts for the ubiquity of certain conceptual oppositions despite cultural change. Very basic conceptual oppositions and associations (like male/female, day/night, or sun/moon) are likely to be found in some form in almost every culture, though each may have different associations. Like other forms of cultural software, their ubiquity can be explained in several ways: These basic oppositions may be dispersed to many different civilizations through memetic descent. Or they may have been independently created because people in many different cultures have faced similar experiences and dealt with them in similar ways.

Conversely, a memetic approach allows us to account, in ways that a Lévi-Straussian cannot, for diversity across cultures and dissensus within cultures. There is no reason that particular conceptual oppositions would be equally salient for all persons in all cultures. For example, we would not expect that oppositions like right/duty or hardware/software would exist in every culture. Over time the kinds of associations that spread widely among members of a culture might change, with older ones transformed and newer ones emerging. Finally, if networks of association are a form of cultural software, they will exist in populations, with slight differences in each person's tools of understanding. Although communication and social learning can harmonize the understandings of a culture's members, no two persons will have exactly the same set of cultural associations.

In this way we can explain convergence of understandings without having to assume Lévi-Strauss's universal structures of the unconscious. And we can explain the forms of divergence, dissensus, and historical change that his theory could not account for. In short, the theory of cultural software shows how we can reinterpret structuralism in terms of individual thought and belief. Structuralism can be freed from its questionable metaphysics, and its insights can finally be appreciated by even the methodological individualist.

Structuralist Homologies

The central focus of structuralist analysis is conceptual homology. A homology is an association of conceptual oppositions. For example, given oppositions between A and B and between C and D, we might have the conceptual homology "A is to B as C is to D," or A:B::C:D. Although this form of analysis is best known through its application to the culture of primitive societies, I shall offer an example taken from contemporary American law: the contrasting roles of the judge and jury in American legal thought.

An American trial involves both questions of law and questions of fact. In the American legal system, citizen juries are the triers of fact, while the judge is responsible for determining the law to be applied to those facts. At the end of the trial, the judge offers instructions to the jury about the law that they must apply to the facts that they find. Reliance on lay juries to find facts and apply law in many different types of legal controversies is a peculiar feature of American legal culture. In England, for example, the jury has been eliminated in most civil cases (except, interestingly, in cases of defamation). In civil law jurisdictions, the lay jury has never had the importance it has had in America.

Gerald Torres and Donald Brewster have noted that the common understanding of the division of labor between judges and juries in the American legal system is linked to a number of conceptual oppositions: these include not only law versus fact but also reason versus passion.⁶ Thus judges, who construe the law, are associated with reason, while jurors, who must sort out the particulars of the factual situation, judge the credibility of witnesses, and act as the conscience of the community, are associated with passion and emotion. The judge is a learned professional who understands legal doctrine, while the jury is unschooled in the law, having only a "brute" sense of justice. The judge presides over many different cases and issues rulings of law that are entitled to precedental effect. A particular jury sits on only one case and is dissolved after doing its work. It does not create durable rules of law but offers a ruling on only one factual situation; its finding has no general precedental effect (although specific findings of fact can bind the identical parties in future litigation).

This analysis produces the following homologies: judges are to juries as law is to fact, as reason is to passion, as formal justice is to informal justice, as the permanent is to the transient, as justice through application of general rules is to justice in the particular case. The association of these conceptual oppositions is a tool of cultural understanding. Like all other such tools, it both enables and misleads understanding. This network of associations allows people to understand, describe, and make judgments about the respective roles of judges and juries, but it simultaneously produces prejudgments, prejudices, stereotypes, and overgeneralizations about judges and juries. In this way the system of networked associations both assists cultural understanding and creates ideological effects.

The above homology explains a surprising amount about how people conceptualize the respective roles of judge and jury. Because the jury is associated with passion, for example, the jury, and not the judge, needs to be controlled. The emotional and discretionary functions of the jury are necessary to justice. Yet the jury must be disempowered, held back, for its own good. Otherwise, it will be a "runaway" jury (a term that connotes a wild beast—the symbol of dangerous passion). Hence American law contains a set of rules of evidence that are said to "constrain" jury discretion (as one would constrain a wild animal or a capricious child). Evidence is excluded when it would unduly prejudice the jury (the evidence of a rape victim's prior sexual history, for example) or inflame the jury (particularly gruesome evidence of a murder). Failure to abide by these rules can be reversible error. Thus to the previous homology of conceptual oppositions we should add the additional opposition of constraining: requiring constraint.

Our semiotic analysis of the way people think about and talk about judges and juries should not be confused with an assertion or a proof that judges really are more rational than juries or that juries really are more emotional than judges. Nor is it a claim that this is the proper or morally appropriate way of thinking about judges and juries. Rather, we are trying to identify patterns of existing cultural thought about judges and juries, and it is entirely possible that this way of thinking may be misleading or unjust.

Moreover, this analysis does not offer us a series of rules that people consciously follow when thinking about judges and juries. People do not recite the homology "judges are to juries as reason is to passion" to themselves and then apply it. Rather, the homology that we discover in symbolic forms is evidence of a cognitive construction that has already occurred, which the semiotician codifies into a set of organizing rules or principles after the fact. The semiotic analysis shows us the results of a series of generative tools applied to a particular aspect of social life. Because of the network of associations that people use in thinking about judges and juries, they construct what they consider appropriate roles for the social institutions of judges and juries, and their associated social institutions. Thus, it is not surprising that semiotics finds a set of structures in what was produced according to these generative tools. Nor is it surprising that, when asked to defend the institution of judges and juries, or the particular features of this institution, people will respond in terms of the network of oppositional categories that structuralism discovers.⁷

Although homologies are not rules consciously followed, they do seem to have an important hermeneutic function. Such analogies offer people a way of understanding conceptual oppositions and the opposed concepts themselves. The homology A:B::C:D not only links A and C (or B and D); it also helps us to understand the nature of A in terms of C, and the nature of C in terms of C. It also helps us understand conceptual opposites (like C and C in terms of each other. A homology of conceptual oppositions is mutually explanatory. It sheds light on the thing to be explained but also reflects light back onto that which is used to explain.

Mediation, Subcategorization, and Nesting

Although people employ conceptional oppositions to understand the world, the world does not always easily conform to on/off categories. Many situations arise that fall between the poles of existing conceptual oppositions; many situations resemble both sides of a given opposition, depending on how they are described or understood.

In fact, many legal situations mediate between the network of oppositions of judge and jury, law and fact, reason and passion. For example, judges in the American system are sometimes involved in factual issues. If judges can be associated with law rather than fact, how does the conceptual system that we have just discussed comprehend this situation? Often the mediation is treated as a subcategorization of one of the terms of a conceptual opposition, resulting

in a new conceptual opposition. Instead of thinking of judicial fact-finding as a mediation between or mixing of the roles of judges and juries, for example, we might think of it as splitting the category of fact-finding into two categories—facts found by judges and facts found by juries. This produces a new opposition of judicial fact-finding and jury fact-finding. The mediation of oppositional categories, or the subcategorization of one term in a conceptual opposition into two opposed terms, produces a second-level opposition, an opposition within an opposition. This phenomenon is called nesting.⁸

I have noted that people understand new conceptual oppositions in terms of older oppositions in a cultural system. The same phenomenon applies to oppositions created from subcategorization. A second-level opposition replicates the associations that are linked with the first-level opposition. Put symbolically, if A:B::C:D, then a':b'::A:B::C:D. If the judge is associated with reason, law, and restraint of passion, judicial fact-finding will be characterized according to these criteria in opposition to the fact-finding of juries (who are associated with undisciplined passion). That is to say, judicial factfinding is to jury fact-finding as judges are to juries, as reason is to passion, and as law is to fact. Judges engage in a more "reasonable" and disciplined type of fact-finding, or a fact-finding that cultivates reason and restrains passion. Sometimes, for example, the parties waive their rights to a jury trial, and the judge acts as the trier of fact. In such cases, judges often dispense with many of the rules of evidence that they use to shield evidence from juries, on the grounds that unlike juries, they will not be unduly swayed (that is, that they are more reasonable). Moreover, because judges already know the law, they can easily sort out the admissible from the inadmissible evidence.

Even in jury trials, the judge engages in some kinds of fact-finding. Often, for example, the judge has to rule on whether certain evidence is admissible by reference to the facts. The judge's fact-finding in these circumstances is in aid of reason because it helps restrain the jury's unreasoning passion. In addition, a judge in the American system must follow the evidence presented at a jury trial in order to act as a check on jury discretion. If the jury's verdict is too unreasonable given the weight of the evidence, the judge may order a new trial or, in extreme cases, direct a different result notwithstanding the jury's verdict. Finally, the judge may remove a certain question from the jury if the testimony presented by the parties is such that no reasonable jury could come to a contrary conclusion. Thus the judge finds facts in order to police the boundaries of reasonableness. Not surprisingly, such questions are called questions of law in the American system, even though they clearly involve factual inquiry.⁹

The concept of oppositions nested within oppositions is a corollary of a fundamental structuralist tenet: the meaning of a cultural event or artifact comes from its opposition to other cultural events or artifacts—in short, the

context in which it is considered. The identification of judges with reason and lack of prejudice arises only in the context of an opposition with juries. Although judges are associated with reason in comparison with juries, there is great concern in jurisprudential debates over the possibility that judges themselves might be swayed by passion and prejudice. In this case, the reason/ passion opposition is no longer identified with judges as a group versus juries as a group; it is now used to understand a division within the category of judges themselves. The context has changed, and with it, the associations between conceptual oppositions. Nevertheless, the terms of this debate are a variation on the reason/passion homology that we have been studying. Indeed, we might even understand it as a subcategorization or mediation of the opposition between judges and juries. The prejudiced judge is to the nonprejudiced judge as juries are to judges. The prejudiced judge is lawless, and therefore needs constraint to perform her proper function; hence she must submit to the constraints of the larger reason of the law. The law itself, which the judge applies, therefore acts as a constraint on the judge; it ensures that her actions are in accordance with the reason of the law. Thus we have a mediation of the original homology. Judge constrained by law: judge unconstrained by law:: reason: passion. Once again, passion is dangerous and in need of control. Once again, reason is in charge of restraining passion.

Homologies and Hierarchies

We thus see two basic features in the structuralist theory of cultural meaning. The first is the notion of clusters of linked conceptual oppositions: judge/jury, reason/passion, law/fact, restraining/needing restraint. The second is the notion of oppositions within oppositions—the idea that mediating categories or subcategories reproduce prior associations in a new form.

Clusters or networks of oppositions do not merely differentiate situations and things. They also have evaluative significance. A differentiation can be a comparative evaluation or a statement of comparative importance. It can also assert a hierarchy. This hierarchy can be one of comparative value, of existing power or status, or of legitimate power or status. Thus the division between judges and juries may carry with it unspoken assumptions about the appropriate distribution of power and authority between them. That is because we are not neutral about the relative value of reason and passion, their relative importance in our lives, or the relative authority that each should have.

If we prefer reason to passion, then we also are likely to prefer that which is associated with reason to that which is associated with passion. We will think it important that reason be in control of passion, and hence that things associated with reason should be in control of those things associated with passion.

Passion is at its best when it serves its appropriate function and is in its appropriate place. That is when it is subservient to reason or less powerful than reason (as in the case of judicial control of juries), when it resembles reason (as in the case of judicial fact-finding or a jury not swayed by unnecessary emotion), or when it acts to further the goals of reason (as in the case of juries supervised by judicial control).

Conceptual oppositions that form hierarchies of comparative value, status, power, or authority can be associated with other hierarchies of comparative value, status, power, or authority. Thus the hierarchy of judges over juries can be buttressed by analogies to other accepted or prevalent forms of hierarchy or comparative evaluation. Put more generally, the hierarchy of A over B can be supported by analogy to the hierarchy of C over D. Thus the homology A: B:: C: D not only explains or clarifies the nature of A and B but also supports the comparative evaluation between them.

Torres and Brewster, for example, have suggested that the way that people talk about juries is also the way they stereotype women. ¹⁰ Juries are said to be capricious and unpredictable, easily swayed by emotion, and yet, on the other hand, intuitive, perceptive, and merciful; while judges are associated with the "masculine" stereotypes of reason, law, rules, and order. This analysis does not claim that men and women are actually the way that masculine and feminine stereotypes portray them to be. It assumes only that these stereotypes exist and that they occur widely in American culture. Thus, the use of "feminine" metaphors about the jury establishes a connection between the dominance of judges and the cultural hierarchy of patriarchy—the ideology that values men and things associated with men over women and things associated with women. Stereotypes that justify one kind of hierarchy become linked to other hierarchical oppositions and serve to justify them as well. In this way, the associated hierarchical oppositions mutually reinforce each other.

Indeed, patriarchy itself is supported in part by a series of analogies to other oppositions in culture and nature. The hierarchical relation of men to women is explained by, understood by, and justified by the relation of this opposition to other conceptual oppositions. These oppositions, in turn, are linked to still further ones, and so on indefinitely.¹¹

We must be careful not to infer from the above example that the judge/jury relationship directly reproduces or supports patriarchy, or that the judge-jury system is an inherently "male" institution. This misunderstands the nature of the ideological effect. People use hierarchies they are already familiar with in order to explain other hierarchies. A hierarchy is made to seem more natural by analogizing it to another hierarchy that already has some cognitive force. Moreover, two conceptual hierarchies may mutually reinforce each other by being associated. Nevertheless, the effect is not perfectly symmetrical, because

some conceptual oppositions are more basic, or more powerful, or more central to our thought than others are. Thus, in theory, the opposition of judges and juries might reinforce the hierarchy of male and female or the opposition between reason and passion just as the latter two conceptual hierarchies tend to explain, justify, and support it. In practice, however, it is more likely that the power of explanation runs largely in one direction. The hierarchical relations of male and female stereotypes and of reason and passion are probably more deeply rooted in our culture than the relation of judges and juries and therefore offer more support to the reproduction of the hierarchy between judges and juries than that opposition does to either of the other two.

The "feminization" of discourse about the jury is more than merely a description and more than an evaluation. It also assigns roles of comparative authority and power. One way of establishing and justifying the superior power of judges over juries and reason over passion is to link this opposition with other oppositions in which one term has greater power over the other, or with oppositions in which one term should have greater power over the other.

In the standard conception of judges and juries, the jury and passion must be associated with lesser power because passion is not necessarily less powerful than reason in all of its manifestations. Passion can sometimes be more powerful than reason; consider a mob or a violent animal. Thus, the jury is "feminized" in discourse in order to associate it with a subordinated, less powerful version of passion—to avoid the dangerous implications of a passion that could be more powerful than reason. This discourse thus trades upon (or rather assumes) the comparative powerlessness of women in order to establish the subordinate role of passion (or its representative, the jury) in a system of law. (More to the point, it trades on or assumes the patriarchal attitude that it is appropriate that women should be less powerful than men.) Thus when Torres and Brewster claim that the jury is feminized and therefore thought less powerful, they borrow an insight from feminist scholars that feminization of a concept is a way of dominating it.¹²

The association of women with less power than men is a standard patriarchal assumption. Even so, I must stress that the concept of women's power is a problem for patriarchal ideology. It is more correct to say, from the standpoint of patriarchal thought, that whether or not women are less powerful, they must be associated with lesser power. The association of women with weakness or with lower hierarchical status must be established and reinforced and reproduced just as much as the association of lesser power for juries must be established, reinforced, and reproduced in the ideological construction of judges and juries. Thus patriarchal thought has many strategies for handling potential associations between women and power. Powerful women are seen as dangerous and unfeminine. Images of powerful women are generally unflatter-

ing stereotypes of deviance: witches, for example, are women who possess the power of magic. The power that women are permitted within patriarchal conceptions must be limited to specified areas, must never threaten the larger power of men, and must ultimately be subservient to male authority. These ideological strategies are never fully successful; even so, they may still be pervasive.

Once again, the structuralist analysis does not assume that the stereotypes implicated in the homology

masculine: feminine:: more powerful: less powerful

necessarily reflect current social realities about men and women. Rather, the assignment of stereotypical male and female characteristics is one of the ways in which subordination of things associated with the feminine is justified or made to seem natural. This assignment is a source of ideological power. The consistent use of feminine metaphors to describe a nonsexual concept like a jury is not so much evidence of differences between men and women as evidence of an ideological strategy of justifying particular relations of power. In the story of power relations, the one who plays the part of "the girl" is the one who is subordinated and whose power or dangerousness must be eliminated. (Indeed, the identification of woman with "girl"—a child—itself reflects a homology of greater power/lesser power.) Thus as a result of this analogical strategy, passion has its proper place within a system controlled by reason. A properly functioning jury can be emotional only as long as it is assigned the sort of values that patriarchy stereotypically associates with women who "know their place"—sympathy, mercy, intuition, and so on. The feminization of the jury thus solves the problem of keeping the jury in its proper place.

Note, moreover, that such analogies work in the opposite direction as well. The identification of the jury with passion and with feminine stereotypes not only justifies its constraint by the judge and the need to subordinate its passion and rough sense of justice to the judge's reason and knowledge of law; it also establishes the connection between the judge and reason, thus justifying the judge's superior position. Just as the jury is feminized (emasculated) in discourse, the judge is associated with the "male" metaphors of reason, authority, and justice.

The Economy of Oppositional Logic

I have argued that networks of oppositions evaluate by combining conceptual oppositions that are evaluative or hierarchical. In our example, the hierarchy of judges over juries was supported by the hierarchy of reason over passion. This discussion assumed a general preference for reason over passion. Never-

theless, in conceptual oppositions like that of reason and passion, we do not always privilege one term over the other, and hence the two terms do not have a unitary set of hierarchical or evaluative associations. Each opposition is a conceptual tool that is used over and over again in many different contexts; through this usage, it becomes linked to many different oppositions, and it forms many different networks of association. Thus, although conceptual oppositions are used to buttress hierarchies or comparative evaluations, they can do so in many different ways, for each has many different associations connected to its terms, and these associations display the terms of the opposition in many different lights.

There are many good or superior associations connected with reason in its opposition with passion, for example, as well as many bad or inferior associations. There are contexts in which reason is viewed as better or more powerful than passion, of greater authority than passion, or needing to be in control of passion, and other associations where the reverse is true. That is because we do not in every case prefer reason to passion or think that reason should always be in control of passion. In matters of artistic expression, a cold, logical, and unfeeling person may be thought to be at a comparative disadvantage with a person who we say is "in touch with" her emotions. In physical activities, people often report that it is better simply to act according to their feelings than to think consciously about what they are doing. To return to the example of legal controversies, legal decisions must be tempered with qualities like mercy, sympathy, and intuition if the law is truly to be just.

Similarly, in the many different contexts in which the conceptual opposition of mind and body is used, mind is sometimes the favored term, while body is sometimes more favored. There are both good and bad associations with mind in its opposition to the body, as well as good and bad associations with the body in its opposition to the mind. Thus we favor body over mind when we prefer the solid, practical, down-to-earth person to the impractical dreamer, the effete intellectual, or the pedant. Conversely, we may favor mind over body if the body is associated with the dangerous elements of passion, or with earthly troubles, while the mind is associated with reason and the more honored "life of the mind." Under this network of associations, the philistine or savage is unfavorably compared to the person of culture, and the grasping merchant to the dispassionate scholar. It follows that the same conceptual opposition can have quite different meanings depending upon the context of associations combined with it. The savage can be noble or brutish, the person of culture wicked or civilized, depending upon the play of associations. The power of ideology, then, seems to come from its ability to emphasize some of these conflicting associations to the exclusion of others.

Let us call the preferred or privileged term in a conceptual opposition the

dominant term and the other term the subordinate term. Let us call associations of superior authority or superior value of a term in relation to its opposite "associations of superiority" or "superior associations" and associations in which the term has lesser authority or lesser value in its relation to its opposite "associations of inferiority" or "inferior associations." We can thus state our point more generally: The terms in a conceptual opposition have both superior and inferior associations, depending upon their place in different networks of conceptual oppositions. When a conceptual opposition is used in a network of evaluative or hierarchical oppositions, one term is dominant and the other subordinate. The superior associations of the dominant term are emphasized and the inferior associations are deemphasized or suppressed. At the same time, the inferior associations of the subordinate term are manifested and the superior associations are deemphasized or suppressed. Even though the other associations are forgotten in the process of buttressing a particular hierarchy or comparative evaluation, they do not disappear. The deemphasized associations of the dominant and subordinate terms do not vanish. They are simply suppressed or forgotten in the particular associational network. They will emerge again in some other cultural context. They can also be used to attack or to deconstruct the evaluative hierarchy. Thus the networks of association that the structuralist identifies have a certain ideological force. But this force is hardly immune from attack. It can always be undermined by bringing into play an alternative network of associations in which some of the previously suppressed associations are brought to light.

Conceptual oppositions like those of mind and body or reason and passion appear in culture in many different ways that look conflicting and contradictory if they are examined together. Traditional patriarchal attitudes provide a good example of how the same conceptual oppositions are employed in quite different ways. Patriarchy generally privileges the male over the female, although, as I shall discuss in a moment, this general privileging is accompanied by privileging the female in distinct and limited areas of cultural life. The conceptual opposition of mind versus body is used in a patriarchal ideology in two ways—one that connects the superior associations of the mind with the male and one that connects the superior associations of the body with the male. The first way simultaneously connects the inferior associations of body with the female, while the second does so for the inferior associations of mind.

Thus, if we take the homology male: female:: mind: body:: reason: passion, we obtain a familiar set of patriarchal stereotypes: Men are reasonable while women are too emotional. Men can distance themselves from a situation and consider its merits abstractly, while women are too contextual and too caught up in the feelings generated by a situation. Men think, women feel. In each of these stereotypes, the male association is preferred, and thus the as-

sociations of mind are preferred to those of the body. Yet if we now consider the homology male: female:: body: mind:: active: passive, we can produce another set of patriarchal stereotypes in which the male association is still favored, even though it is now conjoined to the body. Thus women cannot do jobs that require the strength of men, men are practical while women are impractical, men act while women worry, men work while women gossip, and so on. Here the associations linked to the body are preferred to those of the mind. Jeanne Schroeder has summed up this phenomenon aptly when she states that in patriarchal thought, a thing is privileged not because it is male, but is called "male" because it is privileged.¹³

One might wonder how patriarchy can operate successfully if it employs a conceptual opposition like mind and body in completely contrary ways. Why doesn't the apparent contradiction of privileging and then deprivileging the same term in different contexts lead to the intellectual incoherence of patriarchal thought and thus destroy its power over our imaginations? This situation is puzzling only if we assume that a way of thinking is powerful because it is coherent or orderly. But as I have argued, cultural software is the product of cumulative conceptual bricolage. It is not a rationally designed structure of conceptual relationships, but a historical jerry-built product. Patriarchy, like the panda's thumb, is designoid rather than designed. It may have the surface appearance of design, but closer inspection reveals its historical and adventitious development.

The network of conceptual oppositions that we identify with patriarchy is the product of repeated applications of a relatively small number of conceptual oppositions in new and different situations. Conceptual oppositions can be used in many different ways and contexts because their terms have many different features and many different associations. Conceptual oppositions, in other words, are like versatile tools that can be used in many different ways for many different tasks, with correspondingly different results, and corresponding advantages and disadvantages. Oppositions like mind and body, male and female, dark and light are employed over and over again to understand various features of the cultural world. As a result they are repeatedly stamped, like a marker, onto many different aspects of human culture, and the products that they explain carry these associations with them as they are in turn used to explain or understand further aspects of culture. In this way, conceptual oppositions like mind and body or male and female become ubiquitous through recursion.

Thus our cultural understanding of gender is the result both of the many different conceptual oppositions used to understand the opposition of male and female, and of the many uses of the opposition of male and female to understand other aspects of culture. This process will evidence no grand design, although in hindsight people may see pattern and order; it may even look as

if culture is the result of a grand patriarchal conspiracy. Yet once we recognize the historical construction of patriarchy, we should be surprised by its relative coherence and consistency, rather than by any lack of these features.

Nested Privileging

We have just seen that conceptual oppositions can be used in conflicting ways to buttress a comparative hierarchy or evaluation. We should now note the converse phenomenon: a network of conceptual oppositions does not privilege the same elements to the same degree in all cases. Consider once again the example of judges and juries. The cultural characterization of judges and juries that we have been examining generally privileges the judge over the jury to the extent that it views juries as ruled by passion and in need of judicial control. Yet this general ideological privileging is not a claim that judges are superior to juries in every respect. Indeed, ideological privileging often works by designating particular areas in which the subordinated term appears to be dominant or privileged in some way. A large-scale ideological privileging of judges over juries contains smaller-scale areas or pockets where ideology attributes superior associations to juries.

The system of cultural beliefs about American judges and juries that we have been examining assumes that juries and judges have their proper roles and their proper places in a scheme of adjudication. Although judges need to supervise juries, juries have specific tasks to perform in the American trial system. The system relies on the jury to resolve conflicts of testimony, to judge the credibility of witnesses, and to apply concepts like sympathy, intuition, common sense, and "rough justice." These tasks are considered inappropriate for the judge unless the parties have agreed to a bench trial. If the parties have not so agreed, it is improper for the judge to usurp the jury's role, even though the judge is in charge of controlling the jury and ensuring that it conforms to reason and law. If a judge improperly takes a factual question away from a jury and substitutes her own judgments of fact, for example, a higher court will usually reverse her. In short, within certain specified boundaries, the jury has a role that it performs better than the judge, or, more accurately, in which it has greater moral, political, or institutional authority to act than the judge.

It is therefore incorrect to say that the privileging of judges over juries is systematic in every area of the legal system. Juries and not judges are the privileged institution in some aspects of the trial process. This privileging is limited, however; it is of a secondary or subsidiary order. Although juries have positive associations in particular areas of the system, these areas of the system are in turn subordinated to the areas where judges are privileged over juries. The subordinated term, in short, is awarded first prize, but only in second-

class contests. The jury is celebrated as long as, and to the extent that, it knows its place and performs the tasks assigned to it under the larger ideological schema. An ideological schema therefore works not only by evaluation but also by distribution—that is, it assigns particular places or situations in which the subordinated term will receive positive evaluations or some limited degree of hierarchical position or authority.

One can see the process of distribution at work in the context of traditional patriarchal stereotypes about men and women. Traditional patriarchy associates the opposition of male and female with the oppositions of the market and the family, or the public world of work and the private world of the home. This reflects the homology male: female:: market: family:: public:: private:: outside:: inside:: marketplace: home. But in traditional patriarchy, women are not devalued in all areas of life. Rather, they are specifically valued for their contributions to the maintenance of the home and the family. Moreover, under traditional stereotypes, men are thought incompetent at, or at least not as skilled as women in, the tasks to which women are assigned—housework and child rearing. According to traditional patriarchal views, married men who become full-time "househusbands" are devalued accordingly. They are subjected to social obloquy because under traditional stereotypes they are regarded as lazy, unambitious, parasitic, or effeminate.

Men, of course, do work in the house even within traditional patriarchal stereotypes. This is reflected in a mediation or subcategorization of the original homology. Thus male: female:: public world of work: home:: men's housework: women's housework. Men can be "handy around the house," but this does not refer to child care or to cleaning and cooking but rather to a type of "housework" that is appropriate for men—tasks like repairing machines or electronic equipment, carpentry, plumbing, mowing the lawn, and so on. (Interestingly, the expression often used for this is work "around the house" rather than work *in* the house, which uncannily reflects the homology of male: female:: public: private:: outside: inside.) Under traditional stereotypes, it is thought unusual or even unfeminine for women to engage in this type of activity, even though such tasks are also clearly "housework" in the most literal sense.

When men engage in traditionally female housework, they are subjected to a "double whammy." This double whammy also involves a form of conceptual mediation or subcategorization, but it works in the opposite way from the example of the "handyman." The idea of being handy around the house is a conceptual mediation that preserves male privilege by conceptualizing certain work as male even though it is associated with the home. On the other hand, the phenomenon of the househusband, who deliberately undertakes traditionally "female" housework, is understood as a threat to traditional patriarchal male roles. The mediation is therefore viewed as a bad example, deviant be-

havior or incompetent performance in two directions. First, househusbands are viewed as working in a sphere inappropriate to their talents. So they are assumed to be incompetent at housework and child rearing, or at least less competent than women are. Second, because they adopt "women's work," they are acting in an "unmanly" fashion. They are assumed to lack manly virtues or to be failures at living up to expectations of male identity. The double whammy is the double penalty for crossing a boundary established by an association of two conceptual oppositions—it involves a penalty assessed from both sides of the association. In this case, it is the belief that a man who does women's work is a comparative failure both at doing this kind of work and at being a man.

The double penalty works equally in the other direction. Women who adopt roles, traits, or attitudes that are assigned to men are subject to criticism in two directions. First, their performance in the male occupation and role is suspect. Second, they are criticized as unfeminine. Moreover, even their traits and attitudes are interpreted in negative terms. This is possible because traits and attitudes have both positive and negative associations. Consider, for example, a woman who is relatively unemotional. This trait violates the homology male: female:: reason: passion. Thus she may be criticized as being unfeminine. But she may also be assigned the negative traits associated with lack of emotion: coldness, dullness, ruthlessness, and secretiveness, whereas a male with a similar countenance might be assigned the positive associations of stability, reasonableness, dependability, and depth.

Nested forms of privilege and evaluation are also present in American race relations. At the beginning of this chapter I noted the homology white: black :: knowledgeable : ignorant. The existence of knowledgable, intelligent blacks presents a problem for this homology, which is solved through a nested privileging. Knowledgeable and intelligent blacks are viewed as more "white" than other blacks. Nevertheless, they are also regarded as somehow less accomplished than the most accomplished whites. Stephen Carter has described a "Best Black" syndrome, under which black professional and educational accomplishments are routinely judged against the standard of other blacks but not against those of whites.¹⁴ A successful black scholar in a particular field thus becomes the best black but not the best scholar. This phenomenon offers black intellectuals some degree of status, but it is an inferior grade of status that ultimately suggests that blacks cannot compete with whites.¹⁵ Blacks are considered successful, but only within a greatly circumscribed arena of competition. Carter argues that this phenomenon is exacerbated by race-conscious affirmative action programs, but it probably existed long before these programs arose.

More generally, black accomplishments and culture are often valued and appreciated in relatively narrow segments of society (sports and music, for ex-

ample) which are subordinated to more "serious" concerns. Association of whiteness with the mind and blackness with the body leads to expectations that blacks will naturally excel in things having to do with the body or with the emotions, whereas whites will excel in things having to do with the mind, with intelligence, and with mental discipline.

This opposition between mind and body is reproduced even in the world of American sports, where black talent is particularly valued. Black athletes are thought to display "natural talent," but white athletes are said to become accomplished because of hard work, emotional maturity, and mental toughness. In this way black accomplishments and black culture are permitted to reign, but only in second-class fiefdoms.

Traditional attitudes toward homosexuality offer excellent examples of the ideological power of conceptual homologies and the operation of the double penalty. Many theorists argue that discrimination against homosexuals is linked to the preservation of traditional gender roles and stereotypes, which are both heterosexual and patriarchal.¹⁶ Masculinity and femininity are defined in terms of attraction to the opposite sex. This produces the homology male: female:: attracted to women: attracted to men: attractive to women: attractive to men. This homology, in turn, is linked with the homology male: female:: manly: feminine :: dominant : subordinate. Homosexuality, and especially male homosexuality, threatens this conceptual order because it undermines the clarity of traditional male and female gender identities and hence the clarity of appropriate male and female social roles, authority, and power. By confusing the network of gender associations, gay men appear both to surrender their masculine privileges and threaten the masculine privileges of other males. In like fashion, lesbians threaten the conceptual order because they refuse their roles as wives and mothers within a traditional heterosexual family.¹⁷

Thus homosexuality mediates between the conceptual network of associations that define femininity and masculinity. But because this mediation is seen as destabilizing, it is interpreted pejoratively and assessed a double penalty. Homosexuality is viewed by heterosexual culture both as a failed case of heterosexuality and as a failed attempt at gender crossing. This is because genders are defined in heterosexual culture in terms of desire for the other gender, yielding the conceptual homology heterosexual male: homosexual male:: masculine: feminine:: better example of masculinity: worse example of masculinity. A classic stereotype of homosexuals, for example, is that they display cross-gender behavior: lesbians are viewed as aggressively masculine, and gay men are viewed as effeminate. Homosexual men are stereotyped as wanting to dress in women's clothing or as seeking "feminine" jobs. Because this is not a man's "true nature" (which is to desire women and be associated with "manly" things) the homosexual man is doubly a failure. He is a failure as a man (because

he does not desire women and attempts to behave like them); and he is a failure as a woman (because he is a man).

As part of the double penalty, homosexual men and women are also assigned the negative stereotypes of the opposite sex. Thus gay men are not only viewed as failures as men but are also described in terms of the negative stereotypes of women—as being bitchy or passive, impulsive or overly emotional. Conversely, lesbians are not only viewed as failures at femininity but are also assigned negative associations of masculinity, like aggressiveness and coldness.

Categories as Nested Oppositions

Structuralist analysis shows us that networks of conceptual oppositions produce ideological effects because hierarchies of value, power, or authority are justified and sustained by their association with other hierarchies. A second type of ideological effect stems from the division of the social world into exclusive and opposed categories that are then opposed hierarchically.

To discuss these ideological effects, I will use deconstruction. Many kinds of critiques have come to be called deconstructive. In this chapter, however, I refer to my own theory of deconstruction, which, as I have argued elsewhere, makes the most sense out of many of the textual techniques employed by Derrida and other deconstructionists.¹⁹ This approach to deconstruction is called the theory of nested oppositions.

A nested opposition is a conceptual opposition in which the two terms bear a relationship of mutual dependence as well as differentiation. Put metaphorically, it is a conceptual opposition where the opposed terms "contain each other." This relation of containment or mutual dependence may take one of several different forms, including similarity to the opposite; being a special case of the opposite; conceptual overlap with the opposite; historical, conceptual, or ontological dependence on the opposite; generation from the opposite; or transformation into the opposite over time.²⁰ These relations do not have a single property in common; rather they have a Wittgensteinian "family resemblance." The idea that conceptual opposites are nested is very old—mystical versions of this insight may be found in pre-Socratic thought in the West and in Taoist philosophy in the East.²¹ The idea of a nested opposition is vividly captured in the symbol of yin and yang, which are traditionally portrayed as tearshapes of opposite colors that seem to grow out of each other. In many portrayals of the yin/yang symbol, a small dot of the opposite color is placed in each figure in order to emphasize the mutual similarity and dependence of opposites.

The theory of nested oppositions can be stated simply: every conceptual opposition can be reinterpreted as some form of nested opposition. To deconstruct a conceptual opposition is to view that opposition as a nested opposition.

It is simultaneously to recognize the similarity and difference between opposites or the mutual dependence and differentiation of opposites. The idea of mutual dependence and differentiation between concepts underlies Derrida's notion of différance.²²

Although the theory of nested oppositions claims that every conceptual opposition can be viewed as a form of nested opposition, it does not assert that all nested oppositions are false oppositions or false dichotomies. The fact that two concepts are thought to be similar or mutually dependent in some context does not mean that they are identical in all contexts. This is a confusion of similarity and identity; it betrays an insensitivity to the many changes in context that are the primary concern of deconstruction.

The analysis can work in the opposite direction: we can also find difference among things normally judged to be similar. Instead of discovering that the terms of a conceptual opposition bear a relation of similarity, we can discover a conceptual opposition among things that were previously judged to be similar. In some contexts of judgment we do not notice any conceptual opposition at all but see only similar things. The theory of nested oppositions reminds us that this similarity can become a difference—and hence a conceptual opposition—if the context of judgment is sufficiently altered. Just as deconstruction does not show that all conceptual oppositions are false dichotomies, it does not show that all similarities are false unities. The latter conclusion would also be insensitive to the subtleties of contextual judgment.²³

Although every conceptual opposition can be seen as a nested opposition, we often do not recognize the nested nature of such oppositions. Thus an important ideological mechanism is the suppression of similarity within conceptual oppositions and the suppression of difference within categories. Indeed, to some degree all thought necessarily suppresses the nested nature of conceptual oppositions. The very act of categorization involves the creation of a conceptual opposition (things inside versus outside the category) and the suppression of its nestedness (differentiation within the category and similarity or dependence across the conceptual boundary). This is yet another example of how an ideological mechanism makes use of the same conceptual tools as so-called ordinary or nonideological thought. Indeed, we might even say that the opposition of ideological and nonideological thinking forms a nested opposition whose nestedness has been forgotten.

Suppression and Projection

Let us apply the concept of a nested opposition to the example of judges and juries. In our previous discussion of structuralist analysis, we noted that the homology of

judge: jury:: reason: passion:: law: fact:: rule of law: iustice in the individual case

characterized the respective roles of judges and juries. We can reveal the nested nature of these oppositions in many ways. One way is to reverse the associations. We can deconstruct the homology of judge: jury:: reason: passion by showing how the judge, the law, and the rule of law itself can be the embodiment of unreasoning passion, prejudice, and partisanship, while the finder of fact, and the purveyor of justice in the individual case, the jury, is actually the embodiment of reason and fairness.

It is obvious that as judges are human, they can sometimes be prejudiced, and the rulings that they make will reflect their biases. The judge can therefore be prejudiced or swayed by passion. More fundamentally, however, the law itself can sometimes embody prejudice or unreason. This may seem strange at first because we normally associate passion and prejudice with volatility and changeableness (consider the homology reason: passion: cool: hot:: stable: volatile). Yet the prejudice or the unreason of the law is a prejudice and an unreason in a relatively fixed medium.

How can the law reflect unreason or prejudice? Sometimes the law is too harsh: it fails to take into account the special circumstances of cases, or it lumps dissimilar cases together mechanically and insensitively. In such a case, the jury may ameliorate the harshness of the law by "finding" facts that, when applied to the law, produce a less harsh result. During the era of industrialization in Britain and America, for example, the courts established tort doctrines that allowed defendants to escape liability for negligence (for example, manufacturing defective products or maintaining unsafe working conditions) if the plaintiff also contributed to her injury in any way. These doctrines had such a draconian effect that eventually juries would often "find" that the plaintiff was not negligent at all—even though the evidence indicated the contrary—and would award her a slightly reduced amount of compensation. The sub rosa manipulation of facts by juries might thus be compared to the Solomonic judgment—the judgment that at first looks like an attack of unreason but actually conceals a deeper rationality.

Indeed, sometimes the law may be not simply clumsy and insensitive but actively unjust—it may be the product of bigotry or prejudice or a denial of human rights. In such cases a jury may refuse to enforce the law as written, a process called jury nullification. One of the most famous trials in American history involved the prosecution during the colonial era of the newspaper editor Peter Zenger for seditious libel—that is, for accusing a government official of improper conduct or otherwise holding the official up to public scorn. Under the existing colonial law, the crime of seditious libel was so harsh that one

could not even defend oneself on the grounds that the matter asserted was true. Over repeated threats by the sitting judge, who was determined to force a conviction, the defense counsel argued to the jury that the concept of free speech required the right to speak the truth without fear of prosecution, notwithstanding the positive law of seditious libel. The jury refused to enforce the law and acquitted Zenger. The origins of free speech doctrine in America thus began with an act of jury nullification. Indeed, Akhil Amar has argued that the right to trial by jury was later placed in the American Bill of Rights as a recognition of the importance of citizen juries as defenders of human rights through jury nullification.²⁴ In these circumstances, then, the jury may seem to be a more reasonable institution than the judge or even the law, and it may even be necessary for the jury to exercise control over the law. The possibility of justified jury nullification or sub rosa manipulation of the facts subverts the homology of reason: passion:: judge: jury:: law: fact:: controlling: needing to be controlled. It does so by offering a new homology in its place: reason: passion :: jury : judge :: fact : law :: controlling : needing to be controlled.²⁵

Our deconstruction suggests an important feature of conceptual homologies. The identification of judges with law and reason and juries with facts and passion is not only an explanation, evaluation, or hierarchization. It is also a suppression and a projection. Associating the judge with reason and things associated with reason (for example, the rule of law) simultaneously suppresses elements of judicial behavior that might be unreasonable or "passionate," and projects them onto the judge's opposite, the jury. In psychoanalysis, the term projection is used to describe the patient who sees in others traits, characteristics, or desires that she does not wish to acknowledge in herself. The homology

judge: jury:: rule of law: justice in the individual case:: reason: passion

involves a different kind of projection. It is a suppression or a forgetting of certain traits or characteristics that might apply to one side of the opposition by assigning them to the other side of opposition. By identifying the judge and the law with reason in opposition to passion, the homology forgets, deemphasizes, or suppresses the passion or the unreason that inheres in judges and in the law itself. We might also call this process a distribution: it distributes reason to one side of the homology and passion to the other.

Suppression and projection help to maintain a hierarchy or comparative evaluation in a conceptual opposition. The reason/passion opposition subordinates the latter term to the former. Passion, in its more powerful and dangerous forms, is a threat to reason; if it is to be useful and nonthreatening to reason, it must be seen as under reason's control. In order to be seen as under control, passion must be externalized and projected onto an Other that can be disciplined, subordinated, or supervised. Control through externalization requires that passion (unreason) must be located outside of the materials of the law and the system of rules. Hence the prejudice that might be located in the judge (or the legislature) is projected instead onto the jury, where it can be supervised and subordinated. This control can take the form of actual legal doctrines and institutions, or it can simply rest in the general sense that it is juries and not judges who present a problem that needs addressing. Without this projection, the nonrational element might be seen as emerging from the judges themselves or the law itself; this would destabilize and delegitimate the view of law as reasoned, ordered, and impersonal; it would undermine the comparative evaluation or hierarchy that supports judges and the law.

The deconstructive analysis of homologies has a curious consequence. We have deconstructed the homology of

judge: jury:: reason: passion:: rule of law: justice in the individual case

by substituting a counterhomology. It follows that this new homology, precisely because it is a homology, will also involve its own form of suppression and projection. Thus our attempted subversion is as deconstructible as the original homology it attempted to subvert. We can deconstruct our deconstruction by noting that the portrait of the wise and rational jury it paints is much too optimistic. Few subjects can arouse the concern and indignation of lawyers and judges more than jury nullification. Judges are particularly incensed by lawyers who dare to argue, as did the defense counsel in the Zenger case, that the jury may disregard the law and substitute its own judgment about what is reasonable. Such an argument, they might insist, strikes at the very heart of the concept of a rule of law; more practically, it also strikes at the heart of the control of judges and lawyers over the lay jury. It is no accident that judges make clear repeatedly in their jury instructions that jury members have sworn an oath to obey the law, and that they must apply it whatever their private reservations about the case. (Note the homology implied here of public law: private opinion :: reason : passion.)

The nullifying jury is the epitome of the "runaway" jury, the jury that takes the law into its own hands. Such a jury is the embodiment of unreasoning passion; it is an invitation to anarchy. If a jury can nullify a verdict to protect human rights, it can also nullify a verdict to satisfy its own racial or religious prejudice. If a jury can find Peter Zenger innocent of seditious libel, it can also nullify a homicide charge against a white racist who murders a black civil rights protester. Only adherence to the judge's instructions, and hence to the rule of law, can constrain the prejudices and passions of the jury. If the original homology projected passion and unreason away from the judge and the law onto the jury, our deconstruction of it projects and distorts in the opposite direction. The jury is now viewed romantically as the guardian of rough justice and com-

mon sense, while the unreasoning and dangerous elements of passion are projected onto legal administrators and the content of the law itself.

Neither the deconstruction of the original homology nor its reassertion captures the whole truth about the relation of judges and juries. Each homology involves a forgetting, a suppression, and a distribution of different sides of an opposition to different elements within the homology. Nor is this an accidental feature of this example. Indeed, I claim that this forgetting, this suppression, and this distribution constitute a commonplace ideological mechanism.

A homology distributes opposed properties or concepts (like reason and passion) between two other opposed properties or concepts (like judge and jury). This distribution is always both a projection and a suppression. The homology A:B::X:Y neglects the Y-ness of A and assigns (distributes) it to B; it forgets the X-ness of B and conveys it to A instead. This process works in reverse as well: the A-ness of Y is forgotten and distributed to X, while the B-ness of X is similarly forgotten and distributed to Y. When one finds linked together a large network of associated oppositions—for example, that of

judge: jury:: law: fact:: reason: passion:: rule of law: individuated justice

and so on—it becomes clear that many different forms of suppression and projection are involved and hence many different types of deconstructions become possible.

Indeed, the possible deconstructions of the homology we have been working with seem virtually limitless. For example, our discussion so far has left unquestioned the assumption that passion is a dangerous thing that must be controlled by reason. We merely argued that the dangers of unchecked passion can be present in judges and in the law itself as much as in juries. Thus we still were contrasting undesirable elements of passion to desirable features of reason. But we might challenge these assumptions about reason and passion. We can do this in several ways.

First, we might note that reason in the law can take undesirable forms mechanical reasoning, rigid formalism, and "logic chopping," for example that need to be checked by desirable forms of passion, qualities like mercy, sympathy, and equity. The development of courts of equity to supplement and supervise courts of law is a historical embodiment of this recognition. (We should note that historically courts of equity became just as rigid and technical in their doctrines as the law courts that they were designed to supplement. This suggests that law is continually in need of equity, even in those institutions and doctrines originally designed to promote it.) We thus might argue that reason has dangerous and deleterious aspects that need to be checked by the virtues of passion. This suggests that the just judge, the one who performs her proper function, does not put aside passion and the qualities associated with it

but rather incorporates them into her judging.²⁶ Similarly, as Justice Oliver Wendell Holmes believed, law gains its ultimate legitimacy and intelligibility from its reflection of the sentiments of the larger community.²⁷ This deconstruction thus subverts the earlier homology by uncovering the good or praiseworthy aspects of passion in the judicial function and in the law and asserting their centrality.

Second, we could distinguish between the idea of reason, which we continue to hold a good and desirable thing, and its concrete historical embodiment in institutions like law. We can argue that what passes for reason in law is sometimes not very reasonable—that it conceals its own forms of unreason, bias, and prejudice. In other words, although reason in the abstract may be desirable, its concrete manifestations often possess significant elements of undesirable passion. Conversely, what people think is merely the exercise of passion and emotion by juries and those entrusted with discretionary decisionmaking often turns out to have sound principles of reason behind it. This type of deconstruction works by driving a wedge between an abstract ideal and the concrete embodiment of the ideal. It also attempts to uncover a kind of suppression and projection. By identifying a concrete phenomenon with an ideal (identifying, for example, an actual legal institution with the idea of reason) people downplay the ways in which the concrete phenomenon fails to live up to the ideal model.

Third, we have so far assumed that reason and passion are separate entities. We can also show that the opposition between them is nested by showing that reason and passion depend upon each other, or that concrete examples of reason and emotion have a great deal in common that goes unnoticed.²⁸

Finally, we can show that the opposition between reason and passion is nested by showing how the two terms produce each other or are transformed into each other over time. The reason embodied in law, for example, becomes a kind of passion (prejudice) unless it is made flexible and is subject to questioning and reconceptualization over time. The same is true of beliefs held for a long period of time without question, or of unquestioned obedience to an authority that represents reasoned inquiry.²⁹

Conversely, the passion embodied in a jury's act of nullification—as in the Zenger trial—may cause people to change their minds and see the greater reason in the free speech position espoused by the jury. We see the beneficial effect of passion on reason in many other situations as well: protest and civil disobedience are often necessary to bring home to the audience the injustice of current conditions and laws. Often a highly charged emotional experience (like witnessing a person who is helpless, injured, or suffering) allows people to "see reason" and to change their minds about matters they had been firmly committed to. Thus an experience that touches us emotionally may melt a congealed prejudice that goes by the name of reason and lead to a more rea-

sonable outlook. These examples suggest both that reason can degenerate into prejudice and that passion or emotion can be the wellspring of a more just reason.

Conceptual oppositions have power over human thought because they facilitate thought; and this facilitation is the source of their ideological effects. Conceptual oppositions divide the world into comprehensible categories, but they also suppress similarities and distribute characteristics onto an opposite or an Other. Networks of oppositions help us understand and evaluate the world, but they also reinforce unjustified hierarchies of value and authority. Conceptual oppositions are both necessary and deceptive, ineluctable and illusive; they are consummate examples of the equivocal nature of our cultural tools.

METAPHOR, METONYMY, AND COGNITIVE MODELS

Earlier I mentioned Clifford Geertz's suggestion that cultural understanding operates through classical rhetorical figures like metaphor and metonymy. Cognitive theorists have taken this insight one step further. They argue that metaphor and metonymy are special cases of cognitive modeling that people use to understand the world. People create more complex cognitive models from existing ones through imaginative extension. This process of imaginative extension also resembles the classical rhetorical tropes of metaphor and metonymy.

Cognitive Models

Like other forms of cultural software, cognitive models like those involved in metaphor and metonymy can be distributed widely through human language. On the other hand, many people share the same cognitive models because they create similar ones independently through their own experiences. George Lakoff and Mark Johnson argue, for example, that the most basic cognitive models derive from our experiences as individuals living within a body. Many familiar metaphors are based on models of bodily movement; metaphors of improvement are often based on forward motion. Like good bricoleurs, human beings use their experiences as embodied individuals as models or image schemas to understand other parts of the world around them; these basic models, in turn, are the building blocks of increasingly intricate and complicated cognitive models. Some of the most basic image schemas are those of objects in a container, source-path-goal, linkage, part and whole, center and periphery, up and down,

and front and back.³ All of these primitive schemas are originally derived from bodily movements and embodied experience.

The idea that cognitive structures emerge from bodily experience has a long history.4 The approach has interesting similarities to Pierre Bourdieu's theory of habitus. A habitus is a set of generative principles of understanding shared by members of a given culture. Like Lakoff and Johnson, Bourdieu argues that conceptual development of the habitus often involves analogies to bodily experiences and bodily movements.⁵ A still earlier anticipation of the idea appears in Giambattista Vico's New Science, where Vico states that a universal principle of etymology is that "words are carried over from bodies and from the properties of bodies to signify the institutions of the mind and spirit."6

Although these thinkers stress the role of the human body in shaping very basic conceptual schemas, it is hardly necessary that every human being individually re-create all of her cognitive schemas through bodily experience. Cognitive models are re-created in others through communication and social learning. The theory of cognitive models assumes a historical development from basic sets of perceptual schemas that have a basis in human morphology. Further developments and modifications of these schemas, however, are culturally relative, so we should not expect to find the same cognitive models in every culture. Cognitive models and schemas, like other forms of cultural software, survive and reproduce with different degrees of success in different ecologies. Hence we should expect different cultures to produce different cognitive models and thus to employ different metaphors and metonymies.

Metaphoric Models

Metaphoric understanding operates according to the formula X is Y. The is is not the is of identity. It connotes the modeling of one thing in terms of another, or, more generally, the mapping of one domain of experience onto another.⁷ An example of such a model is the metaphor "understanding is seeing." This metaphor does not assert the identity of seeing with understanding. Rather, it models the process of understanding on the experience of sight. Other common examples of metaphorical modeling include "time is money," "rational argument is war," and "anger is heat." Like networks of association, metaphorical models are cultural heuristics. They help us understand some things in terms of others that we already understand.

Everyday language offers abundant evidence of metaphorical modeling. Because we model understanding on seeing, we also routinely describe understanding in terms of sight and vision, as is evident in expressions like "I see what you mean," "He needs to make his views clear," or "She saw through his deception." Moreover, we often use multiple models to understand the same phenomenon. For example, there are at least two common metaphorical models for arguments. An argument can be a building ("She constructed an excellent thesis") or a journey ("I don't see where she's going with that line of reasoning"). Similarly, we understand time as money (when we spend or waste time) or as movement (when time passes).

Because A is modeled on B, properties of B are assumed also to apply to A, or A will be understood or described as having corresponding features. Thus metaphorical models have conceptual or logical entailments. If an argument is a building, for example, then it must be supported by foundations just as a building is. If its foundations are weak, then the argument is shaky. On the other hand, if an argument is a journey, its conclusion is the end of the journey. If the argument is deficient, one goes nowhere with it, or one goes astray. If the argument is effective, one reaches the desired conclusion—one goes in the right direction. These conceptual entailments are part of the power that metaphorical models possess as heuristics. When we compare A to B, we see elements within A and their relations to each other because we already understand the elements of B and how they relate to each other. Thus a metaphorical model not only describes but also structures understanding. It not only compares but produces cognitive coherence. This is its signal value as a heuristic. A metaphorical mapping imports and applies a ready-made structure that is already understood.

Metaphorical models can be combined with other metaphorical models, producing increasingly complex structures of metaphorical entailment.¹⁰ An example of a complex metaphorical model is the relation between ideas and minds. It combines two metaphors: "ideas are objects" and "the mind is a container." Taken together these produce the model of "ideas are objects contained in the mind." This metaphorical model has many logical entailments: we "grasp" ideas, we "hold" them in our minds. When we understand, we have the right idea "in" our mind: we "get" it.

This model interacts with others in turn. A common metaphor for communication is sending. Under this model, people communicate by sending ideas contained in linguistic expressions (words) from one place to another. When people communicate, the ideas in one person's mind (contained in her words or expressions) travel to and are received into another person's mind. This model combines a number of metaphors: "ideas are objects," "minds are containers," "linguistic expressions are containers" and "communication is sending." The complex combination of these metaphorical models produces a network of logical entailments. We see this in expressions like "Your reasons came through to me," "I didn't quite catch what you meant," "I wish I could put my ideas into words," or "There's a lot packed into what he is saying."

We can find evidence of cognitive modeling not only in common expressions but also in etymology. Metaphorical extension is a useful way of creating new words. Tracing the roots of a word often reveals the metaphorical models that allowed new meanings to be created out of old ones.¹²

Metaphoric models differ from structuralist homologies in two important ways. First, unlike structuralist homologies, metaphorical models or mappings do not necessarily involve relations of conceptual opposition. In the sentence "They camped at the foot of the mountain," for example, the mountain is understood in terms of the human body. The base of the mountain is compared to a human foot, which is the lowest part of the body and supports the body when it stands. This metaphor maps the relation of foot to body (which is a relation of part to whole) onto a physical object. If we were to try to state this mapping in terms of a homology, we would say that body: foot:: mountain: base of mountain. But the relation between the body and foot is not one of conceptual opposition. It involves at least three relations: (1) part to whole (synecdoche), (2) relative position (lowest part), and (3) function (support).

Second, metaphorical mappings make use of schemas or gestalts that cannot always be reduced to relations between two opposed terms. We compare a mountain to a body in the expression "foot of the mountain." Nevertheless, a body has many parts that bear many different relations to each other. The source-path-goal schema underlying metaphors like "life is a journey" obviously has more than two items. The metaphor "ideas are objects in the mind" (which relies on the metaphor that "the mind is a container") involves a comparison between ideas and objects, and between minds and containers. But this metaphor is made possible by a gestalt or visual model that allows us to employ it—a gestalt that includes the understanding that containers have an inside and an outside, that things can be placed in them and out of them, and so on. This gestalt contains many different features, not merely two.

Ideological Effects of Metaphoric Reasoning

Metaphors produce ideological effects because they are selective accounts of experience. Understanding X in terms of Y emphasizes only some features and discounts others. It organizes our imagination about X in one way rather than another. We model X according to the features and relationships between elements found in Y, although we might have modeled it on a completely different set of elements and relations.¹³ A common way of speaking of rational argument in our culture, for example, is through metaphors of war and combat.14 The metaphorical model is "rational argument is war." We speak of demolishing an opponent's arguments, of marshaling evidence, and so on. Taken together, these military metaphors form a coherent set of mutually reinforcing entailments. If rational argument is war, then the other person in the argument is an opponent and the goal of rational argument is to win the argument. One does this by preparing the best defenses, attacking the other person's weak points, shooting down her arguments, and forcing her to capitulate. The parties contend with each other until one party is unable to continue and either surrenders (by agreeing, which signifies that she has lost the argument) or retreats (by changing the subject). In this way the metaphor "rational argument is war" paints an entire portrait of human relations and appropriate behaviors in rational argument.

The potential for ideological effects from this metaphorical model flows directly from the ways in which the model is partial and selective. Comparing rational argument to war captures certain features of rational argument: that individuals strive to better each other in rational argument, that argument is a test of a certain type of strength, and that the participants regard each other as opponents or adversaries. At the same time, this metaphor directs us away from other possible features. The combat metaphor is not the only possible way to understand rational argument. Consider the metaphor of rational argument as cooperative enterprise designed to achieve some mutually desired goal—for example, truth, justice, or accommodation of interests. The parties work together by sharing different theories and perspectives. When agreement is reached, it is not the defeat or overpowering of one's opponent but the satisfactory attainment of a shared goal.

The combat and the cooperation models of rational argument focus on different features of rational argument. Both explain many of the same features of rational argument, but each does so in a different way. Under the combat model, for example, argument can also lead to truth, but the process has quite different ramifications. One example is the adversary system of justice in Anglo-American legal systems, in which the contest of adversaries in adjudication is designed (ideally) to lead to truth. Another is the familiar justification of free speech in terms of a "marketplace of ideas," in which individuals compete with each other to persuade their audience, just as they compete for market share in the economic marketplace.¹⁵ In the adversary model truth emerges not from cooperative striving for accuracy and validity but from the conflict of opposing stories, and from the partial, self-interested motives of the opposing sides. Moving from the example of law to that of science or the humanities, the adversarial model in the academy suggests that truth best emerges from individuals who seek to increase their reputational capital by promoting their pet theories and demolishing the competing theories of other academics.

As we might expect, this model of striving toward truth by demolishing the claims of opponents suppresses other aspects of rational argument—for ex-

ample, the view that truth might best be approached by sharing different perspectives and by attempting to understand perspectives quite different from one's own. The adversarial view also downplays aspects highlighted by a more cooperative view of reason: that progress is cumulative rather than mutually destructive, and that reasoning builds on other people's insights rather than clearing them away to replace them with one's own. In the specific context of legal disputes, the adversarial model deemphasizes the possibility that dispute resolution might best be served by trying to make the parties understand each other's different points of view, with the goal of reaching a mutually acceptable accommodation of their interests.

Thus we can see two different ways in which metaphoric models produce ideological effects. First, metaphoric models selectively describe a situation, and in so doing help to suppress alternative conceptions. By imagining the world one way, we make it more difficult to imagine it in other ways. As Lakoff and Johnson put it, using such a metaphor, "which allows us to focus only on those aspects of our experience that it highlights, leads us to view the entailments of the metaphor as being true," or natural, or what "goes without saying." Metaphors like "rational argument is war," they point out, "have the power to define reality . . . through a coherent network of entailments that highlight some features of reality and hide others."16

Second, and perhaps more important, metaphorical description positively produces social reality as much as it suppresses aspects of it. A metaphor like "rational argument is war" defines and "structure[s] ... what we do and how we understand what we are doing when we argue."17 Thus it is important to understand that a metaphor does more than simply allow us to understand the process of argumentation. This metaphorical model also helps to constitute social conventions of argument and hence helps constitute social reality. Metaphors like "rational argument is war" define the parameters of appropriate social conduct. If argument is war rather than cooperation, we are likely to treat the people we argue with differently, and we will expect different treatment as well.

To be a member of a culture that thinks about argument in these terms is precisely what makes it possible for us to win and lose arguments, for the other person to be an adversary or opponent, for one to be able to gain and lose ground in argument, to plan and execute attacks, and so on.¹⁸

The constitutive function of metaphor, in short, does not simply distort reality. Rather, it makes reality; like narrative construction, metaphor has the power to "make itself true" in social practice. The prevalent social metaphor that rational argument is war creates a series of real expectations about intellectual activity and appropriate behavior in intellectual life that one neglects at one's peril. Indeed, the metaphor even shapes the possible modes of its denial.

Suppose, for example, that a person claims to have won an argument with us. To dispute the claim is already to accept elements of the metaphorical scheme. It becomes quite difficult to avoid talking in terms of winning and losing an argument when others insist upon employing this metaphorical model; to join in the fray means that the metaphor has already worked its power. In politics it is often said that one side has successfully "defined the terms of debate" that the other side must follow. If metaphor were merely a convenient way of describing things, it is difficult to imagine how this could be so. But if a set of socially enforceable conventions and expectations are being created and maintained through our discourse, it is much easier to see how such power could be wielded.

Metaphorical models are classic examples of the ambivalent nature of cultural software. They assist understanding in some respects even as they hinder it in others. Their power stems precisely from their ability to empower understanding by shaping and hence limiting it. To counter this power, we must deconstruct the metaphorical model. We must reveal its metaphorical character by demonstrating that the figural mapping is not logically compelled and showing how it suppresses or downplays important features of a situation. Nevertheless, often one cannot demonstrate this without offering a competing metaphorical model. We may not realize how the adversarial model of argument is limited until we think about argument as a cooperative venture. We need a new vantage point from which to see the limitations of our previous vantage point, a vantage point that a contrasting metaphorical model provides.

This shift of heuristics is characteristic of the way cultural software operates. Metaphorical models assist understanding by prefiguring it; without them, understanding may be difficult or even impossible. Thus one often can counter the power of a metaphorical model and its logical entailments only by substituting a competing metaphorical model that can serve as an alternative heuristic. One can attempt to move parties from confrontation to mediation, for example, by redescribing what they are doing as a cooperative venture. One tries to show that the parties actually have a common shared goal (peace, justice, truth, reduction of costs of litigation, and so on) and that it is in the interests of both to reach that goal.

Metonymic Models

A second kind of cognitive model is a metonymic model. In classical rhetoric, a metonymy substitutes one thing for another that it bears some relationship to. For example, we often associate institutions with their geographical loca-

tions. Thus we speak of "Washington" to mean the U.S. government, the "Pentagon" to refer to the Defense Department, and "Hollywood" to refer to the American movie industry. Note that each of these associations is slightly different. The American movie industry was once actually located in Hollywood, although most of it is elsewhere now; the U.S. government exists in many places, but its capital is Washington. Other implicit conceptual relations in metonymies can be between a piece of clothing and the person who wears it ("Who's the suit?"), a part of the body and a person ("We need a fresh face"), an object and the person who uses it ("There's a new gun in town, sheriff"), a person who controls an institution and the institution controlled ("Bush defeated Hussein"), a producer and a product ("Do you want to listen to some Mozart?"), a place and an event that occurred there ("No more Vietnams"), and a part and the whole ("I've got some new wheels").20 The last example has a special name in classical rhetoric—synecdoche. But for purposes of this discussion I shall treat synecdoche as a special case of metonymy.

A metonymic model generally takes the form "B for A." A, the target concept, is understood as B, the metonymic. The relation between A and B is defined by a conceptual schema. This schema may be a gestalt (for example, a body), a familiar form of association (for example, between users and objects they use), a causal relation, or a standard script or narrative. In our first example, the cognitive schema is the understanding that institutions exist in places. In the sentence "I got here by hopping on a train," the implicit cognitive model is a standard scenario of how one travels in a vehicle. The action of embarkation stands for the entire process of traveling.²¹

Like metaphorical models, metonymic models are heuristics. One reason to understand A as B is that B is more salient or easier to remember. For example, it is easier to think about a typical case than a distribution of differing entities. But like metaphorical models, metonymic models can produce ideological effects because the features of B can be confused with those of A. We may confuse effects with causes, symbols with the things they stand for, parts of a social system with the social system itself, individual actors with the institutions they represent, and so on. Some varieties of what Marxist theory calls reification may be metonymic in character because they confuse products of a social system with the system of social relations and social power.

Some of the most important and pervasive ideological effects of metonymy arise from a special kind of metonymic model, in which an instance of a category stands for the entire category. One example of such a metonym is a social stereotype, another is a paragon or exemplary case. These metonymic models are special cases of the many cognitive models employed in human categorization.

Cognitive Models of Categorization

Human categorization is a complicated psychological process. In the past twenty-five years, psychologists and cognitive scientists have discovered that people often categorize in ways that do not conform to commonsense views about categorization and predication. These commonsense views are sometimes called the classical theory of categories. According to this theory, members of each category share common properties that are necessary and sufficient for membership in the category. These defining properties apply equally well to all members. There are no "second-class" members of a category, at least if the common properties they share have been rigorously defined.²²

The classical model of categories can be a theory of human mental operations or a theory of the underlying logical structure of the world. As a theory of human mental operations, it has come under increasing attack. Numerous philosophical, linguistic, and psychological studies have suggested that human intelligence categorizes in several different ways, some of which do not fit the classical conception of categorization.

We can think of Wittgenstein's famous notion of family resemblance as an early critique of the classical model. Wittgenstein noted that the word *game* does not fit the classical description of a category because there is no single feature that all games share. Instead, games have what Wittgenstein called a "family resemblance"; like members of a family, different games have different properties in common but need not share any single property.²³

A systematic challenge to the classical theory of categories arose out of psychological studies conducted in the 1970s by Eleanor Rosch.²⁴ Rosch demonstrated that people do not experience all members of a category as equally good examples of the category. Subjects asked to rate particular examples of a category (like "bird" or "chair") judge certain examples as more representative than others. Most subjects, for example, rated a robin as a more representative example of a bird than an ostrich or a chicken and rated a desk chair as more representative of the category "chair" than a rocking chair, barber chair, or beanbag chair. Rosch called these more representative examples prototypes. Her experiments showed that such prototypical examples gave rise to many unexpected psychological phenomena, which she called prototype effects.

Prototype effects occur when one member of the category, the prototype, displays asymmetrical or hierarchial relationships with other members of the category. The prototype may thus be seen as more representative of the category than other members. Subsequent experiments showed that prototypes displayed other interesting features. When subjects were asked to identify whether an example was a member of a category or not, response times were generally

shorter for prototypical examples. Subjects were more likely to offer prototypes as examples when asked to list or draw representative members of the category. Subjects applied the notion of similarity asymmetrically when comparing more and less representative examples of categories. Americans who thought of the United States as a highly representative example of a country were asked to give similarity rankings for pairs of countries. These subjects thought that Mexico was more similar to the United States than the United States was to Mexico. Finally, subjects also were more willing to generalize new information about a prototype to a less representative example than to infer that new information about a less representative example was also true of the prototype. Thus subjects were more likely to believe that a disease that affected robins on an island would affect ducks than the other way around.²⁵

Usually a representative example of a category has a bundle of different properties; less representative examples share some of these properties but not all of them, and different examples share different groups of features. The result can be analogized to a central example with different links of similarity shooting out from it in different directions to other less representative examples. Hence this type of category is called a radial category.²⁶ Radial categories manifest prototype effects and Wittgensteinian family resemblances.

It is important to distinguish between the claim that categories display prototype effects and a different claim about categories—that categories have fuzzy boundaries. It is difficult to tell, for example, where the concept "tall" begins and ends. In the case of "bird," however, there is no dispute that robins, ostriches, and penguins are all birds, and yet subjects still report one of these examples as more representative than the others. Wittgenstein's famous example of games unfortunately tends to conflate these two different features of categories—fuzziness and the presence of prototype effects. Wittgenstein asserted that the concept "game" is potentially open and does not need to be fixed in advance to be used effectively, and he also asserted that the concept of "game" had no single property in common but only family resemblances. Prototype effects and family resemblances can occur in concepts that have fuzzy boundaries, but they can also exist with respect to examples where there is no dispute about boundaries.

The psychological research of Rosch and her successors seems to show that human categorization employs several different models, some that appear to behave like classical categories, and many others that display prototype effects and family resemblances. It is by no means clear that these nonclassical models all operate in the same way. Prototype effects are precisely what their name implies—effects of cognitive models—and many different kinds of cognitive models can produce them.²⁷ Psychologists and cognitive scientists are still divided over how many and what different kinds of models are involved in human

categorization.²⁸ Nevertheless, even though different cognitive models may create prototype effects, we can refer to prototype effects collectively as examples of metonymic thinking, regardless of the cognitive model that produces them. They are metonymic because, whenever these effects occur, prototypes, exemplars, or subcategories serve as metonyms for the category. Prototypes, exemplars, and subcategories are used either (1) to represent an entire category, (2) to draw inferences about the entire category, or (3) to provide a normative model for the category. Understanding a category in terms of a prototype, exemplar, or subcategory is often a useful aid to understanding and working with the category and its members. Yet like all heuristics, sometimes metonymic thinking can go terribly wrong.

One of the most pervasive examples of metonymic thinking is the tendency to view all human categories as classical categories. This employs the classical category as a metonym-and hence as the normative and descriptive modelfor all forms of human classification. Precisely because this heuristic is so pervasive, it is an inexhaustible source of ideological effects.

In general, metonymic thinking produces ideological effects when nonclassical categories produce prototype effects whose existence and scope are not fully recognized. As a result, prototypes are improperly employed to make inferences and judgments. In these situations reliance on the assumption that all human categorizations conform to classical categories is misplaced.

Here are some examples of the kinds of prototype effects that, in the right circumstances, can produce ideological effects. They generally fall into two categories: prototype effects that involve some form of stereotyping, and prototype effects that involve some form of unspoken norm.

Overgeneralization

People may assume that all members of a category have the same characteristics as prototypes or prototypical examples. From a prototype of woman, for example, people may infer assumptions about all women's behaviors, preferences, and abilities. The classic example of this ideological effect is a social stereotype.²⁹ Stereotypes can be positive or negative, and they can be derived from other forms of cultural software. Stereotypes about men and women, for example, may be produced by social scripts, conceptual homologies, and networks of association. Social stereotypes are often interrelated: the stereotypical man may be viewed as rational and stable, while the stereotypical women is viewed as intuitive and emotional. These symmetrical stereotypes, in turn, can lead to equally symmetrical overgeneralizations and inappropriate inferences.

Prototypes as Indicators of Relative Prevalence

People may view prototypes or prototypical examples as the most common version of the category, so that other members of the category are viewed as rare, unusual, or exceptional cases, even though in fact these "exceptions" may be just as common as the prototypical examples. In the media in the United States, for example, one often finds a prototype of young black males as uneducated, unemployed, highly emotional, likely to be involved with criminals, gangs, or drugs, easily prone to violence, and likely to get in trouble with the law. When one discovers young black men who do not fit this prototype, they are assumed to be rare and exceptional cases. Indeed, such individuals may be assumed to be especially bright, especially hardworking, and so on. Moreover, viewing them as exceptional and special means that the prototype remains unchallenged and may even be reinforced.

Salient Examples

A special kind of stereotyping involves salient examples. This phenomenon is also related to the availability heuristic discussed in Chapter 8. Prototypical examples often are constructed from or identified with familiar, memorable, or salient examples of a phenomenon (as opposed to typical or frequent examples). People then use these salient examples to make judgments about the probability of events or features of an unknown situation, even if salience and probability are not correlated. Amos Tversky and Daniel Kahneman discovered that when a group of people were asked for the probability that an earthquake causing a major flood would occur in California in 1983, they tended to give higher probability estimates than a similar group who were asked to estimate the probability of a major flood in North America in 1983.30 Politicians often use salient examples to suggest inferences that are unwarranted or at the very least controversial. Ronald Reagan's anecdotes about "welfare queens" driving expensive cars to pick up their welfare checks confirmed a set of stereotypes about poor people as undeserving members of minority groups.

Prototypes as Default Characterizations

People use prototypes as "gap fillers"—to flesh out or supply features of unknown or partially known events or members within a category. If we hear that a man is a bachelor but know nothing else about him, we may create a picture of him in which we assume that he eats out often in restaurants, lives in an apartment, spends lots of money on clothes, and so on. We will use the

features borrowed from a commonly employed social stereotype of a bachelor to fill in what is unknown about him. These features may, of course, be widely divergent from the bachelor's actual existence; he may love to cook at home or be a monk, for example.³¹

Asymmetrical Inferences from Prototypes to Nonprototypes

People routinely assume that inferences about prototypical examples apply to all members of the category, but inferences about nonprototypical cases are not routinely assumed to apply either to the category as a whole or to prototypical cases. Under the ideology of white supremacy, blacks are associated with criminality and whites with lawfulness. (This is another example of prototypes that mirror structural homologies.) If the prototype of the category "black" is associated with criminality, then examples of blacks who commit crimes or black politicians who are corrupt tend to produce or reinforce inferences that all blacks are untrustworthy. By contrast, numerous examples of lawabiding black citizens and upright black politicians do not lead to the inference that blacks are generally trustworthy. Conversely, because noncriminality is projected onto whites, no amount of crime by whites or examples of unscrupulous behavior by white politicians will lead to the conclusion that whites are criminals or that white politicians cannot be trusted. Thus District of Columbia Mayor Marion Barry's conviction for cocaine use reflected badly on all black politicians in ways that Richard Nixon's criminal activities did not reflect on all white politicians.

Prototypes as Reference Points

Prototypical examples are more likely to be used as reference points for comparing or understanding phenomena. Thus if Kansas City is seen as a prototypical American community, then its citizens' tastes and preferences are more likely to be seen as representative of tastes and preferences of Americans in communities like New York, New Orleans, and Miami than the other way around. This heuristic partially explains the often-noted phenomenon of the "white norm" or the "male norm" in everyday understandings of the social world. If people judge males and whites to be more representative examples of human beings than females and blacks, white males become a prototype for the category "human being." Hence the experiences, preferences, and understandings of white males will more naturally be viewed as reference points for evaluating the social world. Conversely, the experiences, preferences, and understandings of women and nonwhites will be seen as less representative, special, unusual, quirky, different, or even deviant.

Prototypes as Evaluative Norms

Prototypical examples can also be viewed as the "best example" in an evaluative sense. Instead of a stereotype that represents the most common features of a category, a prototype can be a norm or ideal. A paragon is one example of a normative prototype. George Washington might be viewed as a normative prototype of a politician. A second kind of normative prototype is an ideal type—for example, the economist's conception of a perfectly rational individual with complete information. A third kind of normative prototype involves a set of norms for appropriate behavior. Thus a normative prototype of a father might be someone who spends time with his children and makes sacrifices for their welfare and their betterment. When people employ a normative prototype, members of the category that fail to match the prototype are to that degree viewed as imperfect, deficient, less valuable, or abnormal. Fathers who do not live up to the normative prototype of fatherhood, for example, may be considered poor examples of fathers. Thus there are two different ways that we can say that an individual is not a good example of a category. On the one hand the individual may be atypical, or may lack many of the most common features of the category. On the other hand, the individual may fail to live up to a normative conception of the category. Often people conflate these ways of not being a good example. This confusion may also produce ideological effects.

Although I have discussed these seven effects of prototypes separately, they often overlap or work in tandem. Lakoff offers a good example of the cumulative effects of metonymic reasoning through an analysis of the expression "working mother." The qualifier "working" suggests that working mothers are somehow special and different from the most representative examples of mothers. Yet changes in economic and social conditions have led more American women to work outside the home than ever before. Why then aren't working mothers prototypical examples of motherhood in American culture?

Lakoff argues that the expression "working mother" is based on an implicit cognitive model of motherhood in our culture based on nurturance. Many other models of motherhood are possible, including being a birth mother, a genetic mother, the wife of the child's father, and so on. The nurturance model of motherhood, in turn, yields a prototypical case of nurturance in our culture—the kind of nurturance that the housewife-mother provides for her children—and thus it offers a prototypical example of motherhood, the housewife-mother. Moreover, there is a prototypical concept of work, which is done away from home and does not include the rearing of children. The term "working mother" is defined in contrast with the prototype of the nurturance-

mother and in conformity with the prototype of work as work done away from home. Because the concept "working mother" is informed by these cognitive models—each with its own paradigmatic cases—it signifies more than simply the intersection of working people and mothers. Lakoff offers the example of a birth mother who gives her child up for adoption and then takes a full-time job outside the home. She is working and she is a mother, but most people would not think that she is a working mother.³³

Thus, powerful forms of metonymic reasoning are implicit in everyday concepts like "working mother." Instead of classical categories of work and motherhood, we have cognitive models that stand for these categories. The nurturance model stands for the category of motherhood as a whole, and the representative example of this model, the housewife-mother, stands for the category of all mothers. This is confirmed by the following English sentences:

- 1a. "She's a mother, but she doesn't take very good care of her children."
- 1b. "She's a mother, but she takes very good care of her children."
- 2a. "She's a mother, but she works away from home."
- 2b. "She's a mother, but she doesn't work away from home."³⁴

Sentences 1b and 2b seem surprising or unusual, while sentences 1a and 2a seem more conventional if traditional. That is because "the word but in English is used to mark a situation which is in contrast to some model that serves as a norm."35 The implicit norm is the nurturing housewife mother.

If the nurturance model serves as the most representative model of motherhood, and the housewife-mother serves as the most representative example within this model, we might expect to see one or more of the following prototype effects:

- 1. People view non-housewife-mothers as exceptional or deviant cases. (Prototype as most common example of category.)
- 2. People assume that characteristics, preferences, and other features of housewife-mothers apply to all mothers. (Prototype as "gap filler" or salient example.)
- 3. People assume that housewife-mothers define the norm of appropriate nurturance and motherhood; conversely, they assume that women who do not correspond to this model are deficient or insufficiently nurturing. (Prototype as paragon, ideal type, or normative model.)

In contemporary U.S. society, a large number of women work outside of the home. Changes in society will gradually undermine the prototype as most common example, although it is also possible that people's estimates of the percentage of women in the workforce lag behind actual numbers precisely because of this prototype effect. Nevertheless, these changes in expectations may not change the tendency to view the housewife-mother as the salient example for inferences about mothers. Equally important, they may not change the tendency to view the housewife-mother as a normative model of appropriate nurturing behavior. Prototypes do not have to be the most common examples to function as paragons and ideals.

The structuralist models that we discussed in Chapter 10 would note similar effects produced by different means. The cultural meaning of "working mother" is produced through a network of cultural oppositions: mother : father :: nurturance : work :: family : market :: private : public. These homologies are both descriptive and evaluative; they distinguish the opposed terms and suggest their appropriate authority and normative significance. For example, work outside the home is differentiated from and privileged over nurturance. At the same time, these homologies "distribute": they assign appropriate roles for which each party or each concept is best suited.

The category "working mother" is a mediation or subcategorization of this homology. Working mothers are opposed both to so-called nonworking mothers, who stay at home, and to working fathers. This produces the following homologies:

"nonworking" mother: working mother:: mother: father:: nurturance: work :: family : market;

working mother: working father:: nurturance: work:: family: market.

In a patriarchal system that privileges women only in roles assigned to the "feminine," working mothers are assessed a double penalty. In both of these homologies working mothers are assigned the inferior associations of each term, and the opposite concepts of nonworking mother and working father receive positive and superior associations. The cultural meanings conveyed are that working mothers are less good at the private responsibilities of family and nurturance than nonworking mothers, and they are less good at the public world of work and exchange than working fathers.

These homologies have ideological power because they implicitly demarcate normal, natural, and privileged associations about mothers and fathers, nurturance and outside work. Many so-called nonworking mothers are less nurturing than many working mothers, and many working fathers are less competent and productive. But the network of associations—with its double penalty—suppresses these facts. It makes the undesirable features (nonnurturance, incompetence) of nonworking mothers and working fathers invisible and projects them onto working mothers.

This example shows how different forms of cultural software can produce similar ideological effects. There are two different explanations for this. The first is that conceptual homologies and cognitive models do not really correspond to different tools of understanding. Each of them is simply a model for the same cognitive process, and the two happen to converge to describe the same basic ideological phenomena. Further experience will show which model is the best description of our understanding, or will reveal an even better model. A second possibility is that there are really several different ideological mechanisms and that they tend to be mutually reinforcing in certain situations, though not in others. This would mean that patriarchal or racist attitudes, for example, are overdetermined by many different forms of ideological mechanisms and ideological effects.

My guess is that both these possibilities have a grain of truth in them. We may eventually discover that the structuralist models developed by Lévi-Strauss and others are special cases of the cognitive models that produce metaphor and metonymy; or we may eventually discover that both types of models point obliquely toward a third as yet undiscovered cognitive mechanism that subsumes both.

Nevertheless, I do not think that all of the various ideological mechanisms that I have described are ultimately one and the same. The heuristics and biases that Tversky and Kahneman discovered cannot easily be reduced to the mechanisms of ego defense implicit in Festinger's cognitive dissonance model, nor can either be readily assimilated into the narrative, structuralist, or metaphorical models of cultural understanding.

In the past four chapters, for example, we have seen how racist ideological effects can be produced by dissonance reduction among subordinate groups, by conceptual imperialism among dominant groups, by faulty inferences from prototypes and salient examples, by conceptual homologies that oppose blackness and whiteness, by suppression and projection of superior and inferior associations, by social scripts featuring stock characters and expectations about ethnic groups, and by recurrent cultural narratives about the American "savage war." It is highly unlikely that all of these effects are produced by the same mechanism. It is much more likely that racial attitudes are produced by many cross-cutting forms of human understanding, which, taken together, have deep roots in our tools of cultural understanding and hence possess great power over our imaginations. Racism and sexism are motley and variegated, despite their admitted power in our lives. They are produced by many different kinds of cognitive tools, and these tools have repeatedly been used to create new ones, carrying into each innovation their potential ideological effects. Unjust attitudes about race and gender are woven deeply into the fabric of our thought; and in this weaving more than one stitch and more than one thread have been used. Ideological mechanisms are the result of bricolage and circumstance; their heterogeneity and disorder are the best evidence of their historical emergence.

A theory of ideology must offer some explanation of how the object of its study—whether beliefs, symbolic forms, discourse, or cultural software—has power over what individuals do and how they think. Thus the study of ideology is necessarily also a cratology—a study of power.

It is possible to have a social theory that ascribes no independent power to the understanding. One could argue that all behavior is structured by the economic or material base of society. Most theorists of ideology, however—including those within the Marxist tradition—have assumed that ideology is an important phenomenon precisely because the way that people understand the world causes them to act against their best interests or to behave unjustly to others.

Conversely, it is possible to have a social theory that overestimates the role of ideology in its theory of power. One might try to reduce military, economic, or technological power to ideological power by arguing that such power is exercised by individuals and groups who are themselves simply the product of ideological or discursive forces. But such a reductionist project is too simplistic. Although ways of thinking do have power over individuals, we must recognize that they do so in concert with many other forms of power that exist in society.

The theory of cultural software offers an account of ideological power. It is a theory of the power of understanding, and hence I call it a theory of hermeneutic power. But it is not a complete theory of power, because it focuses on the power over human beings created by their tools of understanding.

In discussing the relation of cultural software to ideological power, I shall use the work of Michel Foucault as my major foil. I do so because his theory of power/knowledge has been particularly influential, especially given the post-

modern turn from theories of ideology to theories of discourse. I shall argue that Foucault's theory of power/knowledge shares many of the same problems as previous theories of ideology, and I shall argue that the theory of cultural software offers a superior theory of power.

The Study of Ideology as the Study of Power

Foucault's theory of power/knowledge rejects the very term ideology; it attempts to change the focus of inquiry to disciplines and practices of power. Foucault identified ideology with what he understood to be the traditional Marxist model. He objected to it on three different grounds. First, the Marxist model was tied to an unhelpful distinction between economic base and ideational superstructure.1 Once this distinction was made, enormous efforts had to be expended in explaining their proper relationship. In contrast, Foucault argued that knowledge was inextricably intertwined with social systems of behavior. Knowledge arose out of disciplines of knowledge, so that there were various "knowledges" produced by social systems and enforced by their conventions of behavior. Thus one did not have to claim that ideas had power over individuals. Power lay in disciplines and practices, and knowledges themselves were just forms of disciplinary practice.

Second, Foucault objected to the Marxist model because it presupposed a subject who was somehow affected (or deluded) by ideology.² Instead, Foucault wanted to insist that there is no deeper, truer, or more authentic nature of subjects that ideology perverts or disguises; rather, to be a subject is to be created by the various disciplines and practices that exist in one's society.³

Third, Foucault believed that the Marxist model necessarily made a distinction between ideology and truth. But Foucault thought the proper focus of study should be the various discourses through which true and false statements can be made; these discourses themselves are neither true nor false.4

Foucault's general critique of ideology is really directed at a particular instantiation of the concept. None of his objections applies to all theories of ideology; in fact, many different versions either agree with or anticipate his claims. Many theories of ideology (like Geertz's and Thompson's), do not depend upon the Marxist distinction between base and superstructure, and many acknowledge the interrelationship between knowledge and social practice. Althusser's theory of ideology anticipates Foucault's claim that the subject is constituted by culture, as does Geertz's theory of ideology as a cultural system and Gadamer's philosophical hermeneutics. Finally, not all theories of ideology make a sharp distinction between ideology and truth; as discussed in Chapter 5, neutral conceptions of ideology by definition do not do so.

Thus, although Foucault's theories of disciplinary practice, discourse, and

power/knowledge purport to replace the concept of ideology, in fact they bear significant resemblances to many different theories of ideology. This is hardly surprising. As noted previously, Gadamer's theory of tradition, Barthes's notion of a semiotic system, Bourdieu's concept of habitus, and Wittgenstein's concept of a language game also bear significant resemblances to the basic concept of ideology, even if they differ among themselves in important respects. My point is not that ideology should be regarded as a master concept, with all others explained in terms of it, but rather that we should not be blinded by the different terminologies that various theorists use. Each of these theories is concerned with how language and culture have power over individuals' thoughts and actions. Thus one may speak loosely of Foucault's theory of ideology, even though Foucault himself would have rejected the word.

Moreover, although Foucault tried to distinguish himself from previous theorists of ideology and especially from Marx, his theories nevertheless face many of the same difficulties that other theories of ideology face. Even when one changes the focus from ideology to discourse, the basic issues underlying the theory of ideology—the evaluative stance of the analyst, the need for a conception of justice, the causes of ideological effects, the explanation of shared beliefs, the question of ontological commitments, and the problem of self-reference—do not vanish. They simply reappear in new guises.

Within the framework discussed in Chapter 5, Foucault's theories offer a neutral or nonevaluative conception. In *The Archaeology of Knowledge* he does not claim that one episteme is better than another, and in *Discipline and Punish* and *The History of Sexuality* he does not assert that one disciplinary practice is superior to any other. He merely attempts to describe these systems of power and their effects on human beings, or, as he so often calls them, "bodies."

In Chapter 6 I argued that neutral conceptions of ideology ultimately cannot maintain their neutrality because one cannot really describe the effects of ways of thinking on people without evaluative and normative judgments. This point applies equally to Foucault's work. His studied neutrality in description is often merely an attempt to show that the conceptual systems and practices of the past cannot be condemned as easily as we might wish and that they throw an unsavory light on our own current systems and practices. By examining the forms of power/knowledge in the past, we can see how systems of power infiltrate our own lives today. Yet this philosophical tu quoque cannot be performed without a normative conception—in particular, a conception of justice.

On the other hand, Foucault's conception of a subject completely constructed by disciplinary practices raises the puzzling question of why one should even care about what happens to individual human beings if their very individuality is the result of social practice. If subjects are simply the intersection of various disciplinary practices, if they are merely the mouthpieces for various forms of struggle directed by nobody in particular, it is hard to see why we should care about them and their fates. It may be true that bodies are manipulated, watched, cut into, inscribed, and tortured, but that is because they are objects in a never-ending game of power. We should no more feel sorry for these bodies and what happens to them than we should feel sorry for cartoon characters whose ways of thinking and behaving have been drawn by a cartoonist. Perhaps more important, even the sympathy we feel for the victims of surveillance, torture, and other disciplinary practices is simply due to the disciplinary practices of morality and sympathy that are the product of our own cultural moment. We are cartoons crying over the fate of other cartoons.

An analogous problem arises with truth. As Foucault argues, truth is something that arises within the particular discursive structure available to us in the culture in which we live. Different "games of truth" emerge at different points in history and these games cannot be compared in terms of superiority or inferiority, adequacy or inadequacy.⁵ Nevertheless, as Charles Taylor points out, a significant portion of Foucault's project involves unmasking—he wishes to reveal the truth of power behind sanctimonious claims about truth, scientific inquiry, and professional rigor.⁶ Yet the very idea of unmasking implies two things: first, it implies a reality behind surface appearances, and second, it implies a notion of truth that is not limited to our particular cultural moment.

Foucault's insistence on unmasking the operations of power leads to the familiar problem of self-reference: When Foucault acts as ideological analyst, he discovers how power is exercised in a society by agents who do not understand these operations of power. Yet as Foucault himself would hardly deny, his own interests and his own researches are effects of power/knowledge. Thus he is in a position similar to the analysands he studies: his thought, as an effect of power/knowledge, may mask the power relations that constitute it. Like the subjects he studies, he may be blinded to the effects of power on his own thought and behavior. If so, his analyses of the operations of power in other societies (as well as his own) may be limited or distorted, or may miss the mark entirely, just as these subjects could not grasp the existence and effect of the power relations that constituted them.

Foucault's view of the construction of subjectivity has been further criticized by Poulantzas and others on the grounds that it leaves no room for resistance. In response, Foucault argued that resistance arises out of the regime of power itself and is never exterior to it. Indeed, Foucault argued, every regime of power creates its own resistances. Resistance is produced simultaneously with power, much as the back side of an object comes into being as soon as it has a front side. As Foucault says, resistances "are inscribed in [relations of power]

as an irreducible opposite." Thus resistance, rather than being impossible in a network or system of power relations, is in fact part and parcel of its existence.

This answer creates more problems than it solves. It seems at odds with Foucault's general hostility to totalizing explanations of power, as well as his insistence that power works on many different conflicting and overlapping levels at once. Despite this, Foucault's solution to the problem of resistance is Parmenidean: there is only the One—the regime of power—and all resistance to the regime is actually part of the regime. Like Parmenides, Foucault faces the problem of explaining change and motion. Parmenides solved this problem by holding that change and motion were illusions. Unfortunately, this is not a solution easily available to Foucault, because his genealogical method is designed to explain change.

In his writings, Foucault emphasizes that many alternative ways of thinking and living have been crushed underfoot by successive regimes of power/knowledge. How this phenomenon is hard to square with his theory of resistance. Do these alternative ways of thinking and living offer resistance to the regimes of power/knowledge? If they are forms of resistance, then they should be preserved and encapsulated in the new system of power—indeed, they would already exist within it. If they are not preserved, we must assume that they are not the sort of resistance that Foucault's theory explains. (In the alternative, perhaps these previous regimes offered no resistance at all.) In either case, Foucault has not really answered his critics. They want to know how combat with or resistance to a regime of power/knowledge is possible.

Just as Foucault's theory of resistance does not show how an older regime can ever resist a newer one, his theory of resistance does not explain how a new regime can ever supplant or subjugate another that resists it. If the new regime of power replaces another, we must assume that it resists the regime of power currently in place; otherwise, it is difficult to see how it could overcome it. But if it offers resistance, it must already be contained in the existing regime. Hence a new regime of power is impossible. We thus have the curious result that although Foucault speaks of "subjugated knowledges" in his essays on power/knowledge, his theory of resistance gives us no explanation of how they were ever subjugated.¹²

One possible reply to these objections is to revert to a form of humanism. The above arguments assume that networks, regimes, or strategies of power can offer resistance to other networks, regimes, or strategies of power, when in fact only people offer resistance to networks of power. But this resurrection of the individual subject as the locus and source of resistance is unavailable to Foucault. Rather, he would have to acknowledge that resistance is a network of relationships, strategies, and behaviors, just as power is. That is why he is

able to hold that the structure of resistance is built into the structure of power.¹³

Parmenides' doctrine of the One was defended by his fellow Eleatic Zeno, who created a series of paradoxes to show that motion and thus change were impossible. In Foucault's case, however, similar paradoxes would be quite unwelcome, for a modern-day Zeno would use Foucault's doctrine of resistance to show that change in a regime of power is impossible. No regime can ever change into another because no regime can ever be faced with resistance exterior to it that could transform it or overthrow it.

In his archaeological period Foucault claimed that one episteme miraculously transformed into another almost overnight. Given the problems that flow from his theory of resistance, the necessity of making this amazing claim seems more understandable. Foucault could not have asserted that one episteme changed gradually into another, because this would mean that resistance to the episteme could arise from outside the episteme. In any case, Foucault offered his theory of resistance not during his archaeological period but during his genealogical period, when he had seemingly abandoned the structural coherence and totalization of archaeological explanations. Nevertheless, the totalizing character of this theory of resistance is more consistent with the spirit of his earlier work.

In fact, Foucault's theory of resistance is inconsistent with his genealogical approach. His theory of genealogy argues that changes occur from a collision of contending forces or strategies. These conflicts emerge like armies that suddenly find themselves facing each other in a clearing and are thrown into battle. Foucault's theory of resistance, on the other hand, is premised on a closed system in which the structure of resistance is already contained in power relationships. But a genealogical system cannot be such a closed system. It must be the result of evolutionary mechanisms activated by chance events and unexpected changes, invasions, modifications and intrusions. Thus if Foucault is to retain his genealogical approach, he must give up his Parmenidean theory of resistance.

A final difficulty with Foucault's theory of power and knowledge concerns the cognitive mechanisms of knowledge involved in relations of power. Because Foucault identifies knowledge and power with practices, he directs all of his attention to outward manifestations of culture. To be sure, Foucault is not a behaviorist, and he does not appear to have a behaviorist theory of the mind. The problem is rather that Foucault does not seem to have any theory of internal mental processes or cognitive structures. Thus there is a sense in which Foucault is not simply antihumanist, he is also "anticognitivist." First, he pays little attention to mechanisms of cognition and understanding. Second, he writes as if operations of power work exclusively through practices and disci-

plines that are applied to human bodies. Thus he writes as if power arises out of behavior and activity rather than from mechanisms of cognition.

Two features of Foucault's work reflect and reinforce this comparative neglect. The first is his claim that discourses and disciplinary practices constitute the subject. Because these discourses and practices are embodied in technology, symbolic forms, and external behaviors, Foucault does not ask how the mind is in fact constructed by them. He simply takes for granted that mechanisms of socialization and cognition supply whatever is necessary for disciplines of power/knowledge to have their requisite effects.

Nevertheless, disciplines and practices cannot have these effects unless they are understood and internalized by individuals with a cognitive apparatus. Social construction on the order that Foucault proposes requires elaborate mechanisms of understanding that must perform a great deal of work in shaping and constituting the individual's identity and thought. Foucault's account lacks any description or concern with these internal cognitive processes. This criticism should not be confused with a claim that Foucault denies that subjectivity is constituted by culture, or by processes of shared meanings, or even by language. He advocates all of these things. Nevertheless, he wants to advocate all of them without reference to how each individual processes information, or to what is going on inside her head. His view of culture is largely external—it consists in symbolic forms, statements, technologies, architectures, and behaviors.

Foucault's recurrent use of the image of the "body" to refer to human beings also reflects his relative lack of interest in internal mental states. He often speaks of disciplines of the body and of things being done to the body.¹⁵ Foucault's "body" is a metonym for a human being. His use of the term has the obvious rhetorical effect of depersonalizing and defamiliarizing human interaction. But it also has the ideological effects of a metonymic model, as discussed in Chapter 11.

The metonym "the body for the human being" identifies the whole with the part. A metonymic model "B for A" understands A in terms of B and hence may confuse properties of B with those of A. Simultaneously, it suppresses relevant differences between the two. For Foucault, disciplinary power is what happens to bodies: how they are cataloged, inscribed, separated, or gathered together in time and space, how regimens of behavior are prescribed for them, and so on. The problem with this metonym is that bodies can not understand, internalize, or carry out social practices. Bodies do not practice disciplines or devise strategies for dealing with other bodies. In short, bodies cannot understand and cannot act meaningfully, although human beings can. Bodies can be the objects of power/knowledge, but they cannot be its agents—they can be acted upon, but they cannot perform the necessary meaningful actions that sustain a regime of power. Somebody has to be doing something to all of these

bodies. The question is who. The answer is a human being with a particular cognitive apparatus, with historically generated tools of understanding. Expressed in terms of our catalogue of ideological effects, Foucault's emphasis on discourse and practice projects the study of cultural understanding outside of mental processes and onto behaviors, symbols, and cultural artifacts. This leaves only a body that is subjected to these external influences, and a system of power that is "intentional but nonsubjective." ¹⁶

Thus Foucault's recurrent metonym of the body unwittingly symbolizes one the most serious problems of his theory of power. The fundamental difference between understanding human beings as bodies and understanding them as human beings concerns their cognitive processes—their ability to understand and process information and their ability to engage in behavior that has and is understood to have meaning. Foucault's regimes of power/knowledge cannot get off the ground without the cognitive apparatus that makes human beings more than bodies. His theory of power/knowledge thus lacks an account of how the understanding exercises power over the subject.

Foucault's theory also faces the problem of ontology that I described in Chapter 1. For Foucault, epistemes, disciplines, and practices serve the same function as an Objective Spirit or a collective consciousness. They exist over and above individual human beings and are the source of power over them. We might even think of them as "material" versions of these well-worn concepts. Foucault has rid himself of idealism, but he has simply re-created the same Hegelian or Durkheimian formula at a behavioral or material level.

Finally, Foucault faces a problem of differentiation. Foucault's reliance on epistemes, disciplines, and practices explains and enforces uniformity, but at the price of suppressing and failing to explain individual differences in understanding and behavior. Differences among individual understanding and individual behavior are left unexplained or are ignored within the mode of explanation that Foucault offers. Thus his is the most puzzling of genealogies: it is a genealogy without individual variation, which is, in evolutionary explanation, the engine of change.

An important shift in Foucault's work occurs between the first and second volumes of *The History of Sexuality*. In the introduction to the second volume, Foucault suddenly announces that he is shifting course. This change can best be summarized as a movement from disciplines of the body to technologies of the self. The change is not merely terminological. It is in many respects a fundamental reorientation. Rather than asking how social processes arose that did things to human bodies, Foucault now asks how subjects came to understand themselves as selves—how they developed techniques for understanding their place in the social order and the principles of proper conduct within it.¹⁷ This is the question of how an individual comes to recognize his (in volumes

2 and 3 it is almost always a man) ethical duties toward others and appropriate sexual relationships with others.

The sudden shift between volumes 1 and 2 of *The History of Sexuality*, so late in Foucault's career, dramatizes what had been missing in his theory of ideology: the need to understand the processes of ideology from the self's point of view. Thus, Foucault argues, he must now study "the games of truth by which human beings came to see themselves as desiring individuals." ¹⁸

In the final two volumes of *The History of Sexuality*, Foucault rediscovers the individual subject long buried beneath his concern with disciplinary practices. Yet his shift to a "hermeneutics of the self" makes all the more urgent the need to understand ideological phenomena in cognitive as well as behavioral terms.¹⁹

The problems that I have identified with Foucault's theory of power are interrelated. His account of the creation of subjectivity through disciplines and practices must be supplemented by an account of how power arises through the development of cognitive processes and tools of understanding. He needs an account of how power is created through the act of individual understanding, and how this power produces both intersubjective convergence and individual variation in understanding. In other words, to offer a theory of ideological power, one needs to solve the problems of ontology, causation, and differentiation that I posed in Chapter 1. Because Foucault has no satisfactory solution to these problems, his theory of power is also incomplete.

Cultural Software and Power/Knowledge

The theory of cultural software outlined in the preceding chapters is concerned with precisely these matters. It tries to explain the power exercised over individuals because and to the extent that they employ various tools of understanding to get about the world, interact with others, and express their values. Because this form of power arises from the operations of understanding itself, I call it hermeneutic power.

Hermeneutic power should not be confused with the many other forms of power (and forms of violence) that exist in a given society. Hermeneutic power causes us to feel the force of cultural symbols and codes and to behave in accordance with these codes. It bears important relations to other kinds of social power, and other kinds of social power make use of it. But it is not identical with them.

Because his focus is outward to the world of behaviors and practices, Foucault's theory of power/knowledge lacks an adequate account of hermeneutic power. The theory of cultural software can provide such an account. This theory has much in common with Foucault's concept of power/knowledge. First, it offers a genealogical account of cultural development. Second, it asserts that ideological power arises from the nature of subjectivity; it argues that power is implicated in the very acts of knowing and understanding. Third, it holds that power arises from relationships of communication. Fourth, it emphasizes that hermeneutic power is ubiquitous.

The development of the tools of understanding through conceptual bricolage is consistent with Foucault's notion of genealogy. That is because Foucault's "genealogy" is simply another version of the fundamental insight of the philosophy of culture: that much of human culture is the product of the unintended consequences of human action. Although Foucault identifies the idea of genealogy with Nietzsche, his application of it also owes much to Lévi-Strauss's concept of bricolage. In fact, a theory of genealogy is really a theory of bricolage, because it assumes that existing features of culture will be put to new and unintended uses, with new and unexpected effects and developments. Thus when Foucault claims that the genealogist discovers "the secret that [things] have no essence or that their essence was fabricated in a piecemeal fashion from alien forms," he is describing the process of conceptual bricolage in a different way.²⁰

Like Foucault and Althusser, the theory of cultural software argues that ideological power is created by and exercised through the formation of individual subjectivity. The power of cultural software is the power that software has over a person who is partially constituted by that software, who is the person she is because of the software that she possesses. Thus cultural software has power over people because it constitutes people; it produces ideological effects in society because people must make use of it in order to act in society.

The hermeneutic power involved in ideological effects is simply a special case of the power that cultural software has over individuals generally. In essence, the entire previous discussion of ideological effects has been concerned with the mechanisms of ideological power. In previous chapters we have seen how many of the basic tools of cultural understanding that we inevitably and necessarily employ in our understanding of the social world—heuristics, narratives, metaphors, categories, and networks of conceptual associations—shape our thoughts and hence our actions in important ways. Whenever we offer an account of an ideological mechanism, we also explain how it produces power over our imaginations. Thus, within the theory of cultural software, the connections between understanding and power—between ideology and cratology—are fundamental.

Ideological power is an inevitable consequence of the operations of subjectivity, because hermeneutic power is an inevitable consequence of being a person existing in a culture at a particular moment in history. Because individuals must understand the social world through use of their cultural software, they

are inevitably subjected to various forms of hermeneutic power merely by existing as persons equipped with and constituted by cultural software. Each act of cultural understanding is a potential source of ideological power over the individual because each act of understanding is a source of hermeneutic power over the individual. Hermeneutic power, and hence ideological power, is not something wholly imposed on a subject from without; it results from the interaction of the social world with a subject already programmed to receive information in a certain way. As Stanley Fish notes, the force of ideology is not an external force, and ideological power does not operate like a gun at your head. There is no gun at your head: "The gun at your head is your head."

Because individuals are constituted by their cultural software, they are continually immersed in forms of hermeneutic power without noticing it. Thus Foucault's claim about the ubiquity of disciplinary power is also true of the hermeneutic power of cultural software. Take, for example, cultural codes concerning dress. Cultural understandings of appropriate and attractive dress expect women to wear high heels in certain situations. For some, these cultural expectations are oppressive, but they are oppressive in part precisely because they are internalized—the individual feels that she is being forced by community expectations to dress in ways she would rather not. But if a particular individual does not mind wearing high heels and even thinks that they make her look more attractive, she does not feel oppressed or disempowered by the cultural codes that require them.

We may make a partial analogy to the forces of nature. When a swimmer swims with the ocean tide, she does not necessarily feel the tide as a force. Nor do we feel the force of the air that presses against us, unless there is a sudden drop or increase in pressure that produces wind. Nor do we feel the inertial force of the earth's accelerated motion around the sun (produced by a gravitational force), or the solar system's motion within the galaxy. By analogy we might think of hermeneutic power (and ideological power) as a sort of background power that we live within, a power that is constitutive of our everyday existence. Like normal air pressure or the acceleration of the earth around the sun, it is a necessary albeit unnoticed element of our lives, a background force that accompanies and produces our life on Earth. We do not feel the force of the various background forms of ideological power until we oppose them in certain ways. Then we are like a swimmer who tries to swim against the tide and suddenly feels its strength.

The example of air pressure is important for another reason: not only do we not notice normal air pressure but our bodies are designed to operate correctly only within tolerable deviances from this normality. If air pressure becomes too little or too great, we cannot survive. To continue the analogy, there

may be an important sense in which hermeneutic power is not felt in ordinary circumstances partly because our ability to participate in a culture or a shared set of conventions or expectations requires this power to be present. Without this force, our culture, and our cultural identities, could not long survive. The power of cultural software binds members of a culture together and makes following, participating, and developing cultural conventions possible. The fact that this power can be used for good or for ill does not change the fact of its ubiquity; its capacity for good or bad use is implicit in the ambivalent conception of cultural software.

Cultural Software as an Alternative to Power/Knowledge

Although the theory of cultural software bears many similarities to the Foucauldian approach, it also has important differences. These differences resolve the problems that I have previously identified with the Foucauldian model of power/knowledge.

The theory of cultural software differs from Foucault's theory of power/ knowledge in six ways. First, the theory emphasizes that disciplines and technologies of power cannot come into being or be sustained without cognitive mechanisms of understanding. Second, the theory of cultural software is overtly normative and evaluative; it presupposes an idea of justice that subjects the analyst's own beliefs to scrutiny. Third, it provides a catalogue of cognitive mechanisms involved in the constitution of subjectivity and hermeneutic power. Fourth, because the theory locates the source of hermeneutic power in each person's individualized tools of understanding, it does not need to postulate a version, whether material or otherwise, of a collective consciousness or an Objective Spirit. Fifth, for the same reasons, it does not need to assume that all forms of resistance are already contained within a larger system of power. Sixth, because the theory is premised on an ambivalent rather than a neutral or pejorative conception of ideology, it can acknowledge cultural software both as a source of power over individuals and as a source of individual autonomy. In this way it escapes the excesses of Foucault's antihumanist conception.

Foucault explains power through outward manifestations in behavior and practice. The theory of cultural software emphasizes how processes of understanding produce power; it argues that power arises out of cognitive mechanisms as well as out of technology and social practices. In the first chapter I noted that without cultural software our technology becomes useless and our institutions fall apart. One can make the same point about Foucault's disciplines. Disciplines require the existence of cultural software to support and make meaningful the practices and techniques of normalization. Technologies of power require and presuppose cultural software.

Unlike Foucault's histories, the ideological analysis of cultural software does not purport to be nonevaluative. It rests on an ambivalent conception of ideology. As noted in Chapters 5 and 6, this conception necessarily depends upon an idea of justice that can be used to critique the analyst's own views as well as those of the analysand. As we have seen, Foucault's theory of power/knowledge purports to be neutral, but it is really pejorative. The normative bite of his analysis stems from the horror at watching our lives completely constituted by ever tightening chains of power. Foucault's theory has all of the problems that attend nonevaluative theories of ideology: he cannot describe what he wants to describe without taking some sort of normative stand about what is true and what is just.²²

Mechanisms of Hermeneutic Power

Unlike Foucault and his emphasis on external practices and disciplines, the theory of cultural software offers a series of mechanisms that describe how subjectivity is shaped and constituted and how acts of understanding exercise power over the individual imagination. Hermeneutic power over our subjectivity occurs in four ways.

First, cultural software has power over individuals simply because it enables understanding. Enabling and limiting are two sides of the same coin. To enable understanding is always to enable it in certain ways rather than in others. This empowerment opens up certain possibilities for conception and understanding while foreclosing others, in the same way that biological evolution creates possibilities for morphological development by foreclosing others. For example, the historical development of animal structures meant that locomotion would occur through the development of legs but not through the development of wheels.²³ The development of reason through history is the development of certain mental structures, but not all structures can coexist simultaneously. Some ways of thinking may not be possible given the tools available at a particular time. Hence the development of cultural software always directs thought in some ways rather than others; it always makes some kinds of understanding easier than others, and it makes still others impossible given the tools that lie to hand.

Not only is the enabling of understanding a kind of limitation, but understanding itself presupposes a certain kind of structure and hence a certain kind of limitation. As Gadamer points out, in order to understand, we need preconceptions and prejudgments.²⁴ These preconceptions and prejudgments not only affect our understanding, they undergird it; without them, we cannot understand anything at all. Stanley Fish puts it succinctly: an open mind is an empty mind.²⁵ To understand, one needs tools, but these tools are necessarily

better designed for some tasks than for others, just as an automobile is better employed for driving than for brain surgery. Thus cultural software enables by disabling: it opens possibilities for understanding by foreclosing others; it expands our minds by limiting them; it manufactures judgment through partiality; it creates personal freedom through mental regulation; it produces the possibility of insight from a necessary blindness.

Second, cultural software has power over individuals because we come to depend upon it, not only for getting about the world but for our very identities as individuals. Cultural software not only allows us to understand but in doing so helps produce the "we" who understands.

In some respects, the power of cultural software is similar to the power that all tools have over those who regularly employ them and hence come to rely on them. We might offer a partial analogy to our increasing dependence on technology. Our technological tools have a certain power over us because they allow us to do things that we could not otherwise do without them. Because they enable us, we come to depend more and more on them. We use them to perform the tasks of everyday life—indeed, we define the meaning of "everyday life" and our expectations of the normal and the ordinary increasingly in terms of what our technology allows us to do. In this way our technology becomes woven into the fabric of everyday expectations and everyday existence. We fully recognize the power that our technology has over our lives only when it breaks down or malfunctions. A stalled car, a power outage, a crashed computer, or a dead telephone line bring forcefully home how greatly our lives assume and depend upon the existence and availability of certain forms of technology.

Nevertheless, the power and effect of cultural software over subjectivity is, if anything, even more profound. Persons are constituted by their cultural software in a way in which they are not constituted by their cars, their telephones, their bank accounts, or their Xerox machines. One's cultural software cannot be cast aside as easily as one can sell a car, break a Cuisinart, or lose money in the stock market. The tools of understanding cannot be discarded at will. As we noted earlier, even when we attempt to be unbiased or to engage in critical self-inquiry, we are not really discarding our tools of understanding; rather, we are using some of them to think about the adequacy of others, or about themselves.

To be sure, there is an important sense in which personhood includes one's property and one's uses of technology; technology does help constitute us as the people we are. We might even expand our definition of "person" to include a person's possessions and access to technology. But technology does not (yet) seem as fundamentally constitutive of personhood as does cultural software. We use technology instrumentally to further our ends, but our use of cultural

software is more than instrumental, for the person who uses cultural software is partly the thing she uses. Instrumentality usually presupposes a person who employs an instrument. But this person does not come into being until she is constituted by her cultural software. Thus we might say that cultural software is also preinstrumental, in that it creates the conditions for what could constitute an instrumental use of cultural software as well as technology.

Third, power over the subject occurs through the act of change in our understanding that occurs through understanding itself. To understand, we must process information, and this means that we must open ourselves up to the possibility of new experiences and the influence of other persons. To understand is to be susceptible to learning, or—less benignly expressed—to reprogramming. One cannot avoid this possibility; it is a precondition of understanding. To risk understanding is to risk change through understanding, and there is no guarantee that the change will not in some cases be for the worse.

Fourth, the process of understanding through cultural software can have power over us even without significantly altering our cultural software. Ideological power can also arise from the manipulation of our existing cultural software that occurs when we understand others. To risk understanding is to risk not only change but also manipulation. The most obvious example is advertising. Sometimes advertising attempts to forge new associations (Pepsi-Cola with being young and having fun, for example), but at other times it merely exploits the associations that we already have. In the latter case, it does not so much rewrite our cultural software as pander to it. It attempts to "push our buttons"—to invoke powerful images and associations that we already possess in order to cause us to act in certain ways. But advertising is only the most extreme and visible example of a general phenomenon. Communication and the understanding of communication always presuppose the possibility of manipulation. Indeed, what we pejoratively call manipulation is only a special case of a general feature of communication and understanding. Symbols and rhetoric always make use of an audience's cultural software—the common associations, heuristics, and metaphors of individual understanding within a culture to persuade or otherwise affect behavior.

Obviously, the line between the third and fourth forms of hermeneutic power—between changing our cultural software and manipulating it—is hardly clear-cut. Indeed, we would not expect it to be. A person's new tools of understanding must be made out of her old ones. If we assume that all of our communicative experiences have some effect, however slight, on our hermeneutic apparatus, then the line between change and manipulation of our cultural software may be one of degree rather than kind.

The Economy of Hermeneutic Power

Because the theory of cultural software locates the source of hermeneutic power in each person's individualized tools of understanding, it does not need the sort of problematic explanation of resistance that Foucault provided. In particular, it does not need to assume that all forms of resistance are already contained in a larger system of power.

The theory of cultural software does not view individuals as the products of networks of hermeneutic power but rather understands networks of power as the result of interactions between individuals with similar (though not identical) cultural software. This software in turn is continually written and rewritten through these interactions. As described in Chapters 2, 3, and 4, this process creates an economy of exchange and development that regulates similarity of understandings while also producing variation and differentiation. Just as we saw earlier that a Zeitgeist or a "spirit of the age" is an effect of the economy of cultural software, so, too, are the networks of hermeneutic power that exist at any time in society. Because the source of hermeneutic power over the individual lies in each person's individualized cultural software, we do not need to dissolve the subject into some larger set of forces in order to explain social power. Nor do we have to postulate some version, whether material or otherwise, of an Objective Spirit, a collective consciousness, or an episteme that ensures that common social understandings are shared and enforced.

We should distinguish this picture from Foucault's claim in volume 1 of *The History of Sexuality* that power exists at various macro and micro levels that are constantly interacting.²⁶ Foucault's division of different kinds of power is not a reassertion of the individual subject. Rather, Foucault argues that various strategies, disciplines, and practices can exist at larger and smaller levels, in more general and more local spaces, and that the larger and more general forms of power may opportunistically take advantage of the smaller and more local forms in place, even as they can also be said to be produced by them. The interaction that he contemplates is an interaction not between human beings within an economy of power but between different levels or forms of power that produce human beings. Human beings as autonomous agents still do not exist in his system.

Antihumanist accounts like Foucault's are attractive because they appear to simplify the process of guaranteeing intersubjective agreement and shared understandings. Antihumanism responds to an underlying fear in social theory: the fear that the human mind is too private, too closed off, too inaccessible to other minds to explain the shared features of our existence. On the other hand, if culture creates individuals rather than the other way around, if the individual is just the intersection of cultural forces, the problem of intersubjectivity van-

ishes. Yet it is only replaced by new difficulties. The problem now becomes how to explain differentiation and disagreement, or, in Foucault's terms, resistance.

This difficulty highlights the comparative advantages of the theory of cultural software. We are all participants in the economy of cultural software. Each of us is continually engaged in writing and rewriting our own cultural software and the cultural software of others. Thus the problem of accounting for dissimilarity and resistance, which so troubled Foucault's theory of power/knowledge, does not arise. Resistance is guaranteed by the fact that each of us is an individual with unique cultural software. This economy of exchange simultaneously produces similarities and proliferates differences.

Individuals exist in fields of hermeneutic power, continually absorbing and sending out information to one another, continually exchanging memes that travel from one host to the next like viruses, mutating and developing as they spread. Unpalatable as this metaphor may be, it nevertheless emphasizes how limited the privacy of our thinking process really is. The concerns that motivated antihumanism turn out to be exaggerated; for in society we discover not a group of individual minds isolated from each other but a network of nodes of memetic transfer and cultural communication through which information continually flows. We find people constantly connected to others directly or indirectly, constantly assailed by messages, idea-programs, instructions, and signs, constantly exposed to a host of attempts to rewrite or otherwise manipulate their cultural software. The memes in our minds are continually being invaded by memes from elsewhere, and they have constructed elaborate defenses to deal with this constant assault. The amount of hermeneutic power, like the air pressure around us, is enormous, even if it is largely unfelt.

Our inherent susceptibility to change and intersubjective regulation comes from the same features of human existence that have made us susceptible to ever new forms of memetic invasion. As noted in Chapter 3, our ability to absorb informational symbionts may have had distinct evolutionary advantages. Once memes began to spread and take over human minds, however, they paved the way for the absorption and spread of ever new varieties. Precisely because people became good at internalizing memes that might help them, it also became possible for human beings to absorb memes that were neutral or even harmful to their emerging interests as persons. Thousands of years of human civilization have not altered this basic predicament.

There is an interesting connection between this memetic account of our susceptibility and Gadamer's ontological hermeneutics. Gadamer argues that in order to understand we must open ourselves to the possibility of change in our own beliefs.²⁷ One can reinterpret this account of hermeneutic openness in terms of hermeneutic power. Understanding requires openness to the object

of understanding, which requires employment of cultural software to absorb the information we find in the object. As a result, understanding involves susceptibility to the resulting effects produced on our cultural software. Openness is vulnerability, and vulnerability is susceptibility to hermeneutic power. Conversely, susceptibility and vulnerability are preconditions of cultural understanding. To risk understanding is to risk a certain kind of memetic invasion, which can sometimes reconfigure and readjust the memes that we have already internalized.

Thus the economy of cultural software operates effectively because understanding itself is a source of power over the individual. We ordinarily think of understanding in terms of mastery and hence control. Yet this mastery is also a form of vulnerability. For example, there is an old saying that the study of law sharpens the mind by narrowing it. There is much truth in this: legal education does change individuals who enter law schools; they gain knowledge and mastery over certain skills, but at the same time their ways of thinking are altered. They submit to a certain form of reprogramming as the price of their mastery and control.

Because understanding requires the possibility of changes in ourselves, it can be transformative even as it produces new skills and new forms of knowledge. The process of understanding is invasive in the deepest way, for it offers the possibility that we will become different from what we are now through our acts of understanding. The converse is also the case: people may resist understanding precisely to avoid change. The theory of cognitive dissonance argues that people sometimes try not to understand things because the new information threatens their sense of themselves. Information and new experience can change the self, and by changing it, disturb it. Dissonance arises when the self senses a threat to its self-conception. We might even think of some varieties of dissonance reduction as ways for the self to fail at understanding as a kind of self-defense. By selectively remembering events and disregarding recalcitrant evidence, the self attempts to resist changes to the self system that might occur if information were accepted and assimilated into the self and its tools of understanding. The phenomenon of dissonance reduction is evidence of the potential power that change through understanding has over the self, just as some astronomers think that the brilliant light of quasars is evidence of an enormous gravitational pull exercised by a black hole.

Our potential for change through understanding is essential to cultural understanding. Moreover, our ability to participate in culture or in shared conventions or expectations requires that we be susceptible to hermeneutic power. First, our cultural software produces the hermeneutic power that binds members of a culture together and makes following, participating, and developing cultural conventions possible. Second, in order for conventions to be shared,

cultural software must be replicated in members of a culture in a way that allows them to coordinate their activities in cultural conventions, whether this coordination turns out to be benign or malignant. Third, in order to reproduce cultural software in others, people must be able to rewrite one another's cultural software through acts of communication. That is how education proceeds. Thus not only is cultural understanding potentially transformative, it is necessarily so—the transformative features of understanding are necessary for the reproduction, growth, and development of culture.

As a shared way of living and thinking, culture is made possible by our ability to assimilate (and thus be changed by) new meme complexes. Just as we can survive only if we ingest foreign substances into our body as food, so, too, our culture can survive only if our cultural software can be rewritten through interaction with others. Both food and culture enter into us; in normal circumstances, we no more notice the invasiveness of cultural understanding than we do the invasiveness of food. All of this changes, however, when people fear an inappropriate influence or a bad effect from their exposure (or the exposure of others) to certain forms of communication. Then they may fear this communication just as they fear exposure to a poison or a carcinogen in the food they have taken into their bodies.

Hermeneutic Subjection as a Source of Freedom

The economy and distribution of cultural software has important consequences for our conception of individual freedom within a culture. Cultural software is both a source of power over individuals and a source of individual autonomy.

There is a significant temptation to move from the insight that human beings are socially constructed to the assertion that they are socially determined. The picture presented here—of an economy and distribution of cultural software—is social constructivist but not social determinist or antihumanist. Individuals manipulate and rewrite each other's cultural software while themselves being affected (and enabled) by their own cultural software. Thus individuals are both the agents and objects of hermeneutic power. This power does not occur "from the top down" but through a continual process of interaction between individuals, or between individuals and the symbolic forms created by other individuals.

Through acts of communication, individuals mutually participate in hermeneutic power. They are both the purveyors and the objects of this power. Those who can manipulate the forms of our social understanding can gain power over us because they can manipulate or alter the very conditions of our understanding of the social world. On the other hand, they can employ this power only because they, too, are subject to it. They can use cultural software

to persuade others only because they themselves can be persuaded; they can manipulate others only because they themselves are potential subjects of manipulation.

Individual autonomy and subjection to hermeneutic power are two sides of the same coin. Autonomy within culture means the ability to articulate one's values and act according to one's desires. But this is done through cultural software; hence it is done using the very means through which one is subjected to hermeneutic power. Hermeneutic power simultaneously facilitates autonomy and subjection.

This conclusion is consistent with an ambivalent conception of ideology. Understanding involves a kind of power over the self, but not all such power is malignant, just as not all communication is manipulation and not all instruction is brainwashing. Some aspects of hermeneutic power are cooperative and beneficial; others are harmful and deleterious. But the difference between the helpful and the harmful, the enabling and the limiting, is not a difference between that which produces hermeneutic power over the individual and that which does not. Rather, both what we call maturation, or mastery, or freedom, or autonomy and what we call delusion or limitation involve the power of cultural software.

Foucault argued that we should not see truth and power as necessary opposites; hence we can be oppressed by socially constructed "games of truth." In contrast, I contend that we should not maintain a false opposition between the freedom of an individual and hermeneutic power over that individual. Ordinarily we assume that an individual lacks autonomy to the extent and to the degree that someone or something has power over her. Hence the power that cultural software exercises over individuals must be a power that denies them autonomy. This line of reasoning seems to lead us ineluctably from social construction to social determinism and antihumanism. Yet hermeneutic power operates differently. Hermeneutic power and autonomy do not constitute a zero-sum game. The ability to decide, to understand, to interact with others, to articulate and express one's values are all hallmarks of individual autonomy. Yet all of these features are developed through cultural software, which is to say that they are developed by being subject to various forms of hermeneutic power. Being a subject of cultural software—which means being subject to various forms of hermeneutic power-creates degrees of freedom. Hence our attitude toward the development of cultural software must be ambivalent rather than negative. To understand is to be given, at one and the same time, new tools of potential understanding and new chains of potential enslavement, and the two are not easily separated.

Foucault also offers a theory of the productive nature of cultural power, but its contours and consequences are quite different. In *The History of Sexuality*,

Foucault argues that power is positive rather than prohibitory. Power does not merely repress subjectivity but actively shapes and produces it. Relations of "bio-power" produce human personality, desire, and preference rather than simply stifling or blocking them. Relations of power create new expectations and hence possibilities of social interaction and social behavior. Thus the human subject is produced by power rather than being merely subject to its prohibitions. Even the human subject's notions of freedom and liberation are products of this disciplinary power.²⁸

Although Foucault says that power is positive and productive, he does not mean that it is positive in the sense of being a good thing or promoting human freedom. Rather, he believes that as forms of discipline and discourse proliferate, they increasingly entangle us more deeply in webs of bio-power. Thus Foucault offers a pejorative conception of the positive production of power. His view is akin to Max Weber's dread of the iron cage of bureaucracy, or Heidegger's concern about the advance of technology.

Foucault offers as an example of his thesis the development of discourses on sexuality. These discourses did not eliminate sexuality as a concern; quite the contrary, they have made sex increasingly important to us. The secrecy and regulation implicit in sexual repression is merely a tactic in a larger strategy. Sex was made a secret so that we might discuss it constantly, devote enormous energies to examining our sexual motives and urges, and devise, prescribe, and follow different regimens of conduct to ensure that our sexual desires were appropriately and effectively channeled. Hence Foucault says that sexuality is "deployed" as one would deploy a set of forces or supplies for battle. Through the deployment of sexuality, the discourses of sexuality proliferate. As they proliferate, they become a more and more pervasive part of our lives, even in our most determined efforts to keep sex a secret or to regulate it. Indeed, the more effort we put into regulating sexuality, the more important it becomes, and the more we must discuss it and those things related to it.

Even attempts to liberate ourselves from the forms of overt sexual repression are just another method of proliferating sexual discourse. Foucault claims that the notion of liberation presupposes the discovery of a deeper, truer self that is freed to express its real desires. Yet this conception is a sham: the very idea of a deeper truer sexual self is itself the product and the effect of the regime of bio-power. The discourses of sexuality create both the idea of the deeper self and the social apparatus that appears to suppress its "true" nature.

Thus, for Foucault, cultural proliferation is not a means of increasing freedom but rather a means of increasing submission and control over bodies. Even our ideas of liberation are just another ruse, just another opportunity for biopower to infiltrate our lives.

The concept of proliferation is central to Foucault's argument about power.

It has obvious analogies in the world of technology and institutions. Technologies proliferate, because technological developments create new needs and new frustrations and lead to new forms and combinations of technological innovations. Similarly, as Weber recognized, institutions proliferate, creating increasingly complicated institutional frameworks.

Once again, although Foucault's description of proliferation appears to be merely descriptive and nonevaluative, his view of cultural proliferation is essentially pejorative. The result of the proliferation of discourses is the Foucauldian nightmare: an ever tightening network of power exercised over human beings.

Foucault's account of cultural proliferation betrays the deficiencies of a pejorative conception of ideology. His analysis is unidirectional and ignores the problem of self-reference. He gives us no account of how he has been able to recognize and unmask the proliferating devices of power that have fooled everyone else. He cannot explain how he has been liberated to recognize that discourses of liberation are delusory.

In contrast, an ambivalent conception recognizes that the proliferation of culture and cultural tools facilitates and constitutes human autonomy as well as human bondage. An ambivalent conception can explain how Foucault as ideological analyst could comprehend what is happening to him. Among the tools of understanding produced by cultural proliferation are those that allow us (and in particular Foucault) to understand the proliferation of cultural power.

An ambivalent conception does not paint a uniformly rosy picture of culture. It appears optimistic only when contrasted to the Foucauldian nightmare. Once again, an analogy to technological proliferation may prove helpful. We might argue that the proliferation of telephones and the technologies to which they have given rise has thoroughly infiltrated and altered our lives, changed our conceptions of privacy and good manners, and created new ways for us to be harmed (telephone advertisements, anonymous threats, obscene phone calls, and wiretapping, to name a few examples). Moreover, it has subjected us to an ever tightening set of expectations concerning our accessibility to the communicative demands of others. It has produced a need, a desire, and a responsibility to be accessible to others, including supervisors, coworkers, and clients, as well as family and friends. This insatiable demand for accessibility has led to the development of pagers, mobile phones, and cellular units that allow individuals to be in contact with and thus at the beck and call of anyone at any time. Hence we might conclude that the proliferation of this technology has led to an accelerating enslavement of mankind.

This is a pejorative account of technological proliferation. Although accurate in many respects, it is nevertheless incomplete. An ambivalent view of

technological proliferation would note all of these problems. Yet it would also recognize that the development of telephone technology has had definite advantages. It has brought people closer together, saved lives, lowered costs, facilitated the exchange of information, and made possible many of the desirable features of social life that we take for granted today. Unfortunately, the benefits of this proliferation have not come without the costs noted above. Indeed, the two have arisen together, and they are not easily separated. We may attempt to ameliorate these problems through further technological innovations. But these innovations, too, will inevitably produce ripple effects in social structure and social expectations perhaps every bit as serious as the previous innovations had. The only way we can fully eliminate the deleterious effects of our technology is to rid ourselves of it, but then we would have to forswear its advantages as well. The recognition that there is no such thing as a "free lunch" in cultural development—or even in the critique of cultural development—is the essence of the ambivalent conception.

Foucault's concept of cultural proliferation is a one-sided vision of the process of cultural articulation introduced in Chapter 2. There I argued that one of the most important features of culture, and of cultural software in particular, is that it allows us to articulate our values. Because this process is one of bricolage, its adequacies and deficiencies are linked. Thus I argued that an ambivalent conception of cultural development flows from a proper recognition of the process of cultural articulation.

Foucault's concept of proliferation views the process of cultural articulation through the distorting lens of a pejorative conception of ideology. In the first volume of *The History of Sexuality*, Foucault offers an account of the cultural articulation of human values through technology, institutions, and cultural software. What he calls the "deployment of sexuality" is actually a broader phenomenon. It is the cultural articulation of sexual desire, as well as of the various virtues that grow out of this desire or in opposition to it. These include, among other things, our values of love, self-control, propriety, and beauty. Through culture, individuals come to recognize and understand their sexual desires, and they develop conceptions of virtuous behavior with respect to these desires.²⁹

In *The Uses of Pleasure* and *The Care of the Self*, Foucault changes his focus from the proliferation of sexuality to the problematization of sexuality. By "problematization" Foucault means the process of a subject's understanding what constitutes problems or difficulties and why they are considered to be a subject of concern. Foucault's shift of terminology in these last two books mirrors his new emphasis on the subject and the subject's point of view. In fact, problematization, like proliferation, is just another perspective on cultural articulation. It is cultural articulation understood from the perspective of the individual subject. In *The Uses of Pleasure*, for example, Foucault focuses on the

importance in the ancient world of the concept of *sōphrosynē*, or self-mastery. The problematization of sōphrosynē is the flip side of the cultural proliferation of sexual desire. Because sex is constructed and understood in a certain way, the issue of self-control becomes an important value, and vice versa. Moreover, like proliferation, problematization is simply a perspective on the process of cultural articulation of values. Problematization means discovering what is important to us and why it is important. Thus it is a form of cultural articulation of human values.

Foucault's twin concepts of proliferation and problematization are incomplete because they do not recognize that through the articulation of human values, human beings can achieve freedom as well as enslavement. The concept of sōphrosynē provides an example. Through the exercise of self-mastery, the ancients believed, an individual achieved a degree of freedom. Viewed from the standpoint of Foucault's theories, the discourse of self-mastery is just another variety of subjection to bio-power, just like the discourses of sexual liberation in the twentieth century. Indeed, the ancients' conception of self-mastery is a particularly heinous form of enslavement because it is not even recognized as such; instead, it is disguised as a form of freedom.

Yet the ancients were not wholly deluded. Self-control and self-mastery, even though culturally created—and indeed, precisely because they are culturally created—are forms of autonomy. To be able to control oneself—for example, to be able to interact with others without unbridled passion or violence—is a kind of freedom. Similarly, to have knowledge and skill—even if these arise only through cultural practices, and disciplines—is a kind of empowerment.

In his Conjectural Beginning of Human History, Kant uses the story of the expulsion from Eden to make a similar point. For Kant, the fig leaf that Adam and Eve don is the symbol of humanity's entry into culture. Foucault might say that wearing clothes for the sake of modesty marks a beginning of the proliferation of the discourse of sexuality, or the human problematization of sexuality. But Kant sees something more. The fig leaf, he argues, "reflects consciousness of a certain degree of mastery of reason over impulse." In this event, Kant argues, we see "a first hint at the development of man as a moral creature." Finally, Kant notes, the entry into culture marks not only the beginnings of moral rationality but also the beginning of the cultural articulation of values. Through the act of modesty, Kant argues, humanity moves from the experience of sensuous pleasure to an appreciation of the beautiful.³⁰

Thus what Foucault might see as the beginnings of human enslavement to disciplinary practice, Kant sees as the beginnings of the articulation of our values, and the beginnings of the development of our culturally created reason. Moreover, Kant argues, as human beings attain the capacity to reason, they

attain the capacity to be free. Thus both morality and freedom are made possible by the process of cultural articulation.

The most satisfactory approach to the philosophy of culture would temper Kant's optimism with Foucault's pessimism. It would recognize that both thinkers describe the same phenomenon from different perspectives. Through cultural proliferation, human beings acquire new skills, new abilities, and new forms of knowledge; yet in the process they make themselves subject to ever new forms of hermeneutic power. Culture, in short, is a predicament, and the theory of ideology stands as a particularly apt symbol of this predicament. The study of ideology is the study of the deficiencies of our thought, but it is made possible only because our thought has already provided the means to think them. It is the study of the powers exercised over our understanding, but it is accessible only because our understanding has already created the power to understand them. It is the study of the limitations of our imagination, but it is conceivable only because our imagination has already bestowed upon us the freedom to imagine them.

In this book I have tried to explain the phenomenon called ideology and the larger cultural predicaments that give rise to it. I have done so through a master metaphor of cultural software and four subsidiary concepts. Each captures different facets of cultural understanding; each conveys different features of the general argument.

The first concept is tools of understanding. Our tools of understanding enable us to grapple with our world, to understand what is happening in it, to interact with others, and to express and articulate our values. Our tools of understanding are produced through bricolage and recursive manufacture. We modify and reuse the old to create the new. Moreover, each tool, no matter how useful, carries its own limitations, for no tool is perfectly adapted for every occasion. As a result, there are inevitable drawbacks and side effects as our tools of understanding are repeatedly employed for new purposes and inserted into new contexts and situations.

A second concept is the *beuristic*. The tools of understanding are better suited for some purposes than for others, and, one hopes, good enough for the purpose at hand. The notion of a heuristic captures the simultaneous adequacy and inadequacy of our cultural know-how. Each ability created carries with it a necessary disability, each perspective opened up carries with it a necessary blindness. In our cultural software benefit and advantage are yoked together, and so our attitude toward our cultural software can be neither positive nor neutral nor pejorative. It must be ambivalent.

A third concept is the *meme*. It captures the idea of culture as a system of inheritance. Cultural know-how is a product of transmission. It spreads through communication and social learning. It is tied to the past through lines of me-

metic descent. Cultures are populations of relatively similar bodies of cultural software, which survive and reproduce in ecological niches. Each person is a carrier of culture, with a slightly different set of cultural heuristics and tools of understanding. Memes grow, mutate, reproduce, survive, or perish in the ecology of our minds and our technologies of information storage. The evolutionary success of cultural software depends on its ability to spread widely and reproduce itself reliably in a particular ecology.

This leads naturally to the fourth concept, that of the virus. Cultural software is a symbiont, which not only invades the self but also helps constitute it. Cultural know-how is passed from self to self, sometimes with deliberate intention but often without the element of choice. Its evolution is a process distinct from biological evolution that does not necessarily enhance human survival. Cultural software has its own interests in survival and reproduction that may be beneficial, neutral, or even harmful to human interests.

These ideas are brought together under a master metaphor of cultural software. This metaphor emphasizes the role of cultural know-how not only in enabling human thought but in constituting persons as persons. It is the most basic conception because it reflects a most basic feature of human life: we exist as embodiments of cultural information.

Our Informational Existence

All living things embody information in their genetic materials. But what is special about humanity is that we transcend both our genetic materials and our environment. We are more than those creatures for whom our genes formed the original blueprint. And we are more than genetic blueprints shaped by subsequent environmental forces. We are more than a haphazard marriage of nature and environment. We are persons: human beings who embody cultural know-how. Cultural software dwells within us and is part of us.

The human being who absorbs and embodies cultural software, who becomes the incarnation of certain forms of cultural know-how, becomes more than genetic information, more than environmental influence, more even than a combination of the two. We become agents and embodiments of history.

The metaphor of cultural software emphasizes this informational aspect of our existence; not simply the information coded in our genes but the cultural information that is made part of our flesh—that is, incarnated within us. It is encapsulated not only in our thought processes and in the materials of our brain (in ways we do not yet fully understand) but even in our facial expressions, our gestures, and our bodily movements. This enfleshment is best symbolized by the fingers of the jazz pianist, trained not only to respond to the keyboard but to improvise upon it. The pianist's fingers possess a second nature. They

know where to go. But their responses are not foreordained. They are not automatic. The fingers of the pianist respond to the moment, they improvise, they create works of great beauty that never existed and never were thought of before.

Our informational nature is also our historical nature—our being in history. We exist in history and history exists in us. We are imbued with information and understandings peculiar to our time; this information and these understandings will mutate and be passed on to be embodied by still others. We find ourselves in a great chain of historical being; we exist in lines of memetic descent, in which we play roles not fully acknowledged or understood. Alongside the course of human events—the wars and famines, earthquakes and diseases, the rise and fall of mighty empires—there is the transmission of cultural software, multiplying and mutating, culminating and dissipating, dispersing and rejoining. We travel and participate in a vast sea of knowledge, custom, and convention, lifting us up, taking us we know not where.

Billions of years ago, a great tide of genetic information and genetic transmission began, a tide that still carries us and of which we are still an integral part. Only a few million years ago, a new tide arose on this planet—a tide of cultural information and cultural transmission. It has steadily gained power and influence, using us as its partly witting and partly unwitting vehicles. Through human technology and colonization, this tide has reshaped the ecology of our planet, confronting and redirecting the older tide of genetic transmission like two great waves colliding on a rocky shoreline.

The collision of the genetic and cultural tides is not only the result of overpopulation and pollution. Our science has made us conscious of the genetic tide itself and how to manipulate it. We have already learned to shape the genetic information of plants and animals in primitive ways, to limit and extinguish other species by brutal choice and careless accident. Soon we shall be able to reengineer our own genes. Then the two tides of genetic and cultural information will swirl around each other, reshaping each other in ways we can only guess.

The Career of Reason

Human reason is an integral part of the tides of memetic evolution. It has a cultural and historical component. And because it has this component, human reason is not a finished product. It is an ongoing project, a collection of historically accumulated tools of understanding, each imperfect and provisional, which metamorphose and meld, spreading and dissipating throughout human populations. Human reason is a feature of populations and cultures as much as of individuals. We are its carriers and its developers, its subjects and its agents.

Through the evolution of culture, knowledge is made flesh and dwells within us.

Throughout this book I have portrayed the devices of human thought and their historical evolution as the source of both understanding and misunderstanding, of both empowerment and confusion. Some will think that this portrait debases reason, or makes it impossible for reason to improve itself and see through injustice and illusion. It does not. Such misunderstandings reflect, I think, the failure to be fully reconciled to the ramifications of our historical existence.

One might object that the picture of reason as an assortment of ambivalent tools fails to explain the adequacy and efficacy of human reason. Beyond the various devices of human thought must there not lie another, purer faculty of reason, which lacks the ambivalent character of all the others and therefore arbitrates over them all? Perhaps our understanding does make use of metaphors and metonymies, heuristics and narrative structures. But surely our ability to deconstruct them indicates that there is some further general faculty of reason that allows us critically to reflect on them. For if there were no general faculty of "good reasoning," how could we see through the cognitive illusions that this book describes? The various heuristics and devices of thought-including metaphor, metonymy, and narrative construction, among many others—are mere supplements to this purer form of reasoning, invoked when convenient, but ultimately unnecessary to critical reflection. In the alternative these devices merely provide raw materials that this other higher faculty of reason sorts, culls, and purifies without needing to employ them in the process of purification.

This objection rests on two confusions. First, it confuses belief in the existence of better and worse ways of understanding the world with belief in a separate capacity of critical reason that arbitrates over lower and more fallible forms. Second, it wrongly assumes that if human reason is a motley collection of tools of understanding, it cannot be efficacious, self-reflective, and selfcorrecting.

Behind this objection is a familiar desire—a desire to preserve human reason from its imagined detractors. It seeks to preserve the power and purity of human reason by identifying some part of human understanding as "reason" and attempting to separate and distinguish it from the remainder. This strategy projects error and illusion onto this remainder in order to reassert the power and mastery of what it labels "reason." But our processes of understanding cannot be divided and separated in this way.

Human beings can and do discover the better and the worse argument. Metaphors can be deconstructed, analogies can be dismantled, narratives can be dismembered. But we do all of these things using cognitive tools (like language) that in other contexts and situations can have ideological effects. There is no pure analytic capacity of "good reason" that is separate from the many devices of human understanding. Reason is a bundle of devices that build on each other and counteract one another's ideological effects. Good reasoning is not so much a matter of purification as a form of triangularization, of imagination and reconsideration, in which we attempt to make use of the many different tools we possess.

Taken by itself, each of our cognitive tools has weaknesses and limitations, yet taken together each can compensate for the others' respective deficiencies. Human reason is like a collection of slender twigs, which, taken separately, bend and break easily, but when bundled become difficult to snap. Human reason is like a roof made of a motley assortment of overlapping materials, which individually let in the cold and the rain, but woven together provide a relatively effective shield against the elements.

If the mind is the product of evolutionary forces, both natural and cultural, the nature of reason could hardly be otherwise. The human brain arose not as a general-purpose problem-solving machine but as an organ that solved particular evolutionary problems—how to recognize danger, how to find food, how to find a mate, how to engage in social cooperation and punish defectors, how to avoid contagious disease, and so on. Evolution is conservative and economical: It always solves the problems before it, not the more general difficulty that might arise at some point in the future. It always draws on the devices available to it; it does not redesign from scratch. So when Nature designed us to be able to recognize defections from social cooperation, she did not necessarily optimize our abilities at psychological introspection or mathematical calculation. When she enabled us to organize expectations of events in narrative form, she did not necessarily optimize our ability to do analytical philosophy. When she instilled a healthy respect for certain indicia of health, she did not prepare us for an era in which these heuristics and behaviors might be counterproductive. Rather, what biological evolution tends to produce is a collection of special-purpose gadgets that work tolerably well for specific environmental challenges, even if they lack more general abilities and efficiencies.

Along these lines, cognitive scientists recently have suggested that the mind might be fruitfully compared to a sort of Swiss Army knife, containing multiple reasoning capacities called "Darwinian algorithms." If we take an evolutionary approach seriously, we recognize that the mind is as motley as it is powerful. The mind is a collection of tools of understanding, each fairly good at the tasks for which it evolved but relatively limited outside its domain. There is no general-purpose faculty of reasoning and problem solving, but together, the various gadgets that we collectively call reason can do an acceptable job.

Like the bundle of twigs or the thatched roof, the mind's performance is

not flawless. Our minds display interesting gaps in abilities, much as our senses occasionally deceive us in optical illusions. As with optical illusions, we can work around them by using our other faculties. One might think that our ability to work around cognitive illusions supports the notion of a general faculty of "good reasoning." Indeed, it demonstrates precisely the opposite proposition. These cognitive illusions, these gaps and lapses in our cognitive competence, are proof that our reasoning powers are the process of evolutionary bricolage, that we are dealing not with a smooth undifferentiated surface of reason but rather with a mosaic of overlapping materials, the joint product of natural and cultural development. Both the existence of these lapses and our ability to compensate for them are signs of evolution at work.

Indeed, if our faculty of reason were smooth and undifferentiated, if we did possess a general-purpose faculty of reason, this would be a strong argument against our minds' having been the product of evolutionary development, whether natural or cultural. The gaps and inadequacies of our reasoning process are evidence of the evolutionary origins of the mind, both natural and cultural. A mind produced by evolution will display both "spandrels"—abilities that later prove useful but which are mere side effects of previous evolutionary design—and "panda's thumbs"—compromises of design created from previous materials that work tolerably well but imperfectly. The person who demands a general, undifferentiated faculty called "good reasoning" does not understand that she is also asking for a being who is not the product of temporal forces of evolution.

One might fall back on the hope that culture's overlay on our mental faculties has successfully smoothed out its rough edges. After all, we have developed language and propositional argument, science and experimental methods. But cultural tools are also historical products: they are the evolutionary result of generations of memes that were able to take root in human minds and spread widely to the minds of others. The cultural component of reason is also a collection of new gadgets superimposed on and merging with the older ones that we have inherited from previous development. Together, this set of tools can recognize and solve many problems. Together, the tools of our understanding can produce what is roughly equivalent to a general-purpose problemsolving machine. But it still betrays its rough edges, its gaps, its inefficiencies. And it is still limited in many ways.

Even so, it is important to distinguish the claim that reason is motley from the claim that reason is unreflective—that it cannot improve itself through conscious analysis of its own beliefs and operations. Perhaps this is the real source of the objection to the picture of reason that I have offered in this book. If there were no separate capacity of "good reasoning"—for example, one represented by propositional discourse—one might fear that human beings could not rationally reflect on cognitive illusions and improve their thinking processes. They would be doomed forever to be the slaves of unreflective customary modes of thought.

Ironically, this objection is itself an example of bad reasoning, for the conclusion does not follow from the premises. Precisely because human thought is self-reflective, it must have a layered, heterogenous, and cumulative character. Human thought is the product of bricolage, and the modification and improvement of human thought through reflection and argument is part of the process of that bricolage.

Human beings can and do transcend unreflective prejudice and custom. The conflict of human wills creates the occasion for conscious reflection about our factual and normative beliefs: these reflections can be assimilated to become the background assumptions and tools of a later era. Through sociability and strife, through human cooperation and human competition, our cultural software reflects on itself, criticizes itself, and modifies itself. Hence there is a continuous dialectic between custom and reflection on custom, between habitual practices of thought and criticisms of these practices, between what is considered "reasonable" at any point in history and reasoned attacks on this rationality.

If human reason is the product of such a dialectic, we would not expect our reasoning abilities to be smooth and unified. Rather, we would expect them to be jagged and variegated. Human reason would tend to look like an old building in a perpetual state of renovation, with old walls halfway broken down, new plumbing joined to older lines, electrical wires shunted through ancient walls, bits of old plaster peeping through newer layers, and dust and refuse everywhere.

One cannot have it both ways. If human reason is to be improved through reason, it must bear the marks of renovation. It must be cluttered, unkempt, and untidy. It must be improved in some respects and disturbingly recalcitrant in others. And it will always be so, as long as the renovations continue.

Indeed, precisely because human reason is corrigible, always capable of selfimprovement, because it responds and develops in the face of experience, it will always continue to be limited in some ways, better at some tasks than others. This, too, is a consequence of its historical production. Biological evolution does not perfect organisms in the sense that it produces creatures equally well adapted to all environmental challenges. Quite the contrary, it tends to produce creatures exquisitely adapted to the environments they find themselves in. By analogy, we can expect that the forces of cultural evolution will not produce forms of human reason equally good at every task. Our tools of understanding always respond to the problems handed to us; they are devised to solve these problems and not others. We cannot know in advance what all

of these problems will be, even though some of them will surely be the unwitting consequences of our own previous actions. The human mind will not eventually become a general-purpose problem-solving machine because life does not present us with general-purpose problems.

The belief in the ultimate perfection of human reason is a temporally extended version of the belief in the human mind as a general-purpose problemsolving device, which is, in turn, yet another version of the belief in a pure, unsullied form of rationality that arbitrates over all of the other facets of human understanding. There is no such smooth, undifferentiated device, equally good at responding to all of the problems and difficulties that may be thrown at it. All tools, precisely because they are useful, are more useful for some things than for others. This trade-off is inherent in the nature of design, and it does not vanish, even when our tools become more sophisticated. Quite the contrary, for trade-offs of design often become increasingly urgent as technology grows in sophistication.

I noted earlier that we human beings exist in a great tide of informational evolution. Yet our participation in the tide of cultural evolution does not mean that we lack agency. Our cultural software surely affects our behavior; our actions always have unintended consequences. But it is a far cry from recognizing this to inferring that we are mere instruments of memetic evolution. We must reject a simplistic either/or view which insists that either we are in full control of the development of our memes or they are in full control of us.

Human beings are not simply an inert environment in which memes compete and breed. Our minds select and reject, combine and reconfigure the memes we are exposed to. We do this both consciously and unconsciously, both deliberately and as a side effect of everyday life. We are active participants in the growth and spread of cultural software, even if we do not have full control over the terms of its evolution.

Just as we must not confuse our subjection to hermeneutic power with a lack of freedom, we must not confuse our role in the development of cultural software with a lack of agency. Being the bearers of cultural software does not eliminate our agency but, rather, creates it, shapes it, brings it into being.

Human beings imbued with cultural software are unique and remarkable creatures. They are knowlege made flesh, produced through the interaction of their biological capabilities and memetic invasion. Through this interaction they transcend the power of both their genes and their memes. They become agents of culture and, equally important, agents of justice.

Ironically, the only thing beyond our choice is whether we will affect the growth and development of cultural software. For this growth and development is history itself; and we cannot absent ourselves from history, just as we cannot exit from culture. We take part, every day, in the production and reproduction

294 | UNDERSTANDING

of cultural software, through language, through participation in social conventions, and through all of the various systems of social meaning. We can avoid affecting the ecology of memes only by refusing to act, refusing to understand, refusing to innovate, refusing to create, refusing to communicate in any fashion at all. But that would be utterly foreign to our natures. For we are beings made of knowledge, and we must communicate to live. To participate in the growth and development of cultural software is our historical destiny. It is our informational fate.

1. Tools of Understanding

- 1. Ethics of the Fathers, Philip Birnbaum, ed. and trans. (New York: Hebrew Publishing, 1949), 5:8, 40.
- 2. Indeed, from its inception the concept of ideology has always been contested, and hence the theory has generated many variations. Compare the variety of definitions offered in Terry Eagleton, *Ideology: An Introduction* (London: Verso, 1991), 1–2.
- 3. Here I am gathering together what proponents of a discourse model deliberately wish to distinguish among. They focus on acts of speaking, writing, and meaning rather than on beliefs. I have no quarrel with the claim that thought, meaning, language, and action are inextricably related. My point is that a pejorative conception of ideology has a particular interpretive attitude toward the object of its critique, whether that object is belief or discourse.
- 4. Compare P. N. Johnson-Laird, *The Computer and the Mind* (Cambridge: Harvard University Press, 1988); Hilary Putnam, *Representation and Reality* (Cambridge: MIT Press, 1988); and John R. Searle, *Minds, Brains, and Science* (Cambridge: Harvard University Press, 1984).
- 5. I should note that the very attempt to divorce these issues is itself controversial. See Gerald M. Edelman, *Bright Air*, *Brilliant Fire*: *On the Matter of the Mind* (New York: Basic, 1992).
- 6. Howard Gardner, *The Mind's New Science* (New York: Basic, 1979), 41. The same, I am afraid, must be said of much of the most important and valuable work in the philosophy of mind. John Searle is the most notable exception, but of course he has also been highly critical of the computational metaphor. See Searle, *Minds, Brains, and Science*, 28–41. In fact, there is an important connection between his critique of the computer metaphor and his views about the study of culture. Searle has argued that what differentiates the study of the social sciences from the study of the natural sciences is that the products of culture are the products of intentionality, something he claims existing computers do not possess (82–83). Thus, at least from Searle's perspective, it

would not be at all surprising that work employing the computer metaphor tends to bracket away questions of cultural understanding.

- 7. Jerome Bruner, Acts of Meaning (Cambridge: Harvard University Press, 1990), 11.
- 8. For an accessible discussion of brain physiology explaining why such a simplistic hardware/software model must be wrong, see Edelman, Bright Air, Brilliant Fire. Moreover, the fact that human beings exist in bodies is an important feature of how their cognitive tools emerge and develop. See Francisco J. Varela, Evan Thompson, and Eleanor Rosch, The Embodied Mind: Cognitive Science and Human Experience (Cambridge: MIT Press, 1991); Hubert L. Dreyfus, What Computers Can't Do: The Limits of Artificial Intelligence (New York: Harper and Row, rev. ed., 1979), 235-55. The metaphoric and metonymic models described in Chapter 11 are premised on the importance of embodied experience to human cognition.
- 9. See William H. Durham, Coevolution: Genes, Culture, and Human Diversity (Stanford: Stanford University Press, 1991).
- 10. If the theory of ideology is properly part of the philosophy of culture, the philosophy of culture is also the philosophy of history, for it asks how people exist as members of a culture in history.
- 11. I distinguish the ability to speak a particular language from linguistic ability in general. There continues to be considerable debate among linguistic theorists concerning the scope and the parameters of innate linguistic ability.
- 12. Hans-Georg Gadamer, Truth and Method, Garrett Barden and John Cumming, eds. (New York: Crossroad, 1975), 245-53.
- 13. Hans-Georg Gadamer, "The Problem of Historical Consciousness," in Interpretive Social Science, Paul Rabinow and William M. Sullivan, eds. (Berkeley: University of California Press, 1979), 103-59.
- 14. Jürgen Habermas, "A Review of Gadamer's Truth and Method," rpt. in Understanding and Social Inquiry, Fred R. Dallmayr and Thomas A. McCarthy, eds. (Notre Dame: Notre Dame University Press, 1977), 335-63; Hans-Georg Gadamer, "Rhetoric, Hermeneutics, and the Critique of Ideology: Metacritical Comments on Truth and Method," and Jürgen Habermas, "On Hermeneutics' Claim to Universality," rpt. in The Hermeneutics Reader, Kurt Mueller-Vollner, ed. (New York: Continuum, 1992), 274-92, 294-319, respectively.
- 15. Stephen Turner, The Social Theory of Practices (Chicago: University of Chicago Press, 1994), 49.
 - 16. Gadamer, Truth and Method, 245-53, 261-62.
 - 17. Ibid., 351.
 - 18. Turner, The Social Theory of Practices, 44.
- 19. Cf. Gadamer, Truth and Method, 263-64 ("It is enough to say that we understand in a different way, if we understand at all").
- 20. David Lewis, Convention: A Philosophical Study (Cambridge: Harvard University Press, 1969), 56, 78, 118.
- 21. As examples, think of racist attitudes, or the cultural meanings of miniskirts. These examples of shared meanings are a far cry from the classic examples of coordi-

nating conventions like deciding whether to drive on the left-hand side or the righthand side of the road. Ibid., 5-8. Moreover, describing conventions as solving "problems of coordination" puts altogether too rosy a glow on social conventions like slavery, or cultural associations of femininity with submissiveness. As described more fully in Chapter 3, we must try to understand how self-replicating conventions and institutions can be parasitic on the human capacity for sociability and harmful to human interests.

- 22. See Immanuel Kant, Critique of Pure Reason, unabridged ed., Norman Kemp Smith, trans. (New York: St. Martin's, 1929), A 346-47, B 404-5.
- 23. See, e.g., Edmund Husserl, Ideas: General Introduction to Pure Phenomenology, W. R. Boyce Gibson, trans. (New York: Collier, 1931); Claude Lévi-Strauss, The Raw and the Cooked: Introduction to a Science of Mythology, vol. 1, John Weightman and Doreen Weightman, trans. (New York: Octagon, 1970); Noam Chomsky, Reflections on Language (New York: Pantheon, 1975).
- 24. The idea of cultural software differs from the Gadamerian notion of a tradition in yet another way: Cultural software encompasses more than linguistic ability. It includes bodily skills that, although teachable through language, are not the same thing as linguistic ability. These include the ability to cook a soufflé, play a musical instrument, or hit a baseball. Although Gadamer insists on the importance of language as the medium of tradition, his formulation fails to encompass all of the many different types of skills and bodily movements that can be transmitted and reproduced in individuals, that constitute them as individuals, and that affect their understanding of themselves and of the world.
- 25. A hardware/software combination of this type is sometimes called a virtual machine, because it uses the software to imitate another machine that has a different hardware configuration or is dedicated to a different set of tasks. For example, with the right kind of software, a Macintosh computer can become a "virtual" IBM-compatible computer and run some kinds of DOS-based programs.
- 26. For an evolutionary argument describing how the capacity to employ software might have developed in humans, see Daniel C. Dennett, Consciousness Explained (Boston: Little, Brown, 1991), 182-?1. Dennett contends that "software" transforms the hardware of the brain into virtual machines that perform various tasks (211). He then argues that human consciousness is the product of these hardware/software interactions (218).

2. Bricolage and the Construction of Cultural Software

- 1. The claim that cultural software is constitutive of the person is also true, in a somewhat different way, about technology and institutions. Our subjectivity may also depend on our participation in social institutions, and it may even depend, as Hegel argued in his theory of property, on the material objects that we own.
- 2. The most well-known philosophical critique of the homo faber model is Hannah Arendt, The Human Condition (Chicago: University of Chicago Press, 1958).
- 3. Donald A. Norman, Things That Make Us Smart: Defending Human Attributes in the Age of the Machine (Reading, Mass.: Addison-Wesley, 1993); R. L. Gregory, Mind in Science: A History of Explanations in Psychology and Physics (Cambridge: Cambridge Uni-

versity Press, 1981); Daniel C. Dennett, *Darwin's Dangerous Idea: Evolution and the Meanings of Life* (New York: Simon and Schuster, 1995), 377–78.

- 4. See Max Horkheimer and Theodor Adorno, *Dialectic of Enlightenment*, John Cumming, trans. (New York: Continuum, 1994; orig. pub. 1944).
- 5. Conversely, one can critique forms of violence or war to the extent that they objectify individuals and deny them recognition as human beings.
- 6. Michel Foucault, *The History of Sexuality*, vol. 1, *An Introduction* (New York: Vintage, 1980), 26.
- 7. We find an analogous result in cross-cultural classifications. The concept of machismo that was articulated in Hispanic and Latino culture is quite different from the virtue of being a mensch among Eastern European Jews.
- 8. Here I follow the excellent discussion in T. K. Seung, *Intuition and Construction:* The Foundation of Normative Theory (New Haven: Yale University Press, 1993).
- 9. Claude Lévi-Strauss, *The Savage Mind* (Chicago: University of Chicago Press, 1966), 16–36.
- 10. See, e.g., Jean-François Lyotard and Jean-Loup Thébaud, *Just Gaming* (Minneapolis: University of Minnesota Press, 1985).
- 11. The development of the ancestral word for arm into the Latin *articulus* would be an example of metaphorical or analogical extension—from a thing to things similar to it in some respect. The development from *arm* to the homonym *arms* is an example of metonymic extension—from a thing to things associated with it. As I shall explain in more detail in Chapter 11, metaphoric and metonymic extension are important features in the construction of ideological thought.
- 12. Pierre Bourdieu, *The Logic of Practice* (Cambridge: Polity, 1990), 86–97; Pierre Bourdieu, *Outline of a Theory of Practice* (Cambridge: Cambridge University Press, 1977), 109–58.
- 13. See Bourdieu, The Logic of Practice, 250-70; Bourdieu, Outline of a Theory of Practice, 87-95.
- 14. Thus, long before Darwin, Immanuel Kant argued that mankind develops through a process of cultural evolution. Kant, "Idea for a Universal History from a Cosmopolitan Point of View," in *Kant on History*, Lewis White Beck, ed. (New York: Macmillan, 1963). However, Kant's vision of evolution, like those of his contemporaries, and unlike Darwin's, was based on a notion of gradual progress toward enlightenment.
- 15. See Stephen Jay Gould, *The Panda's Thumb: More Reflections in Natural History* (New York: Norton, 1980), 77–84. Darwin actually recognized several mechanisms of evolution, of which natural selection was the most important. See Charles Darwin, *On the Origin of Species by Means of Natural Selection*, in *The Portable Darwin*, Duncan M. Porter and Peter W. Graham, eds. (New York: Penguin, 1993), 111.
- 16. Although natural selection is the central mechanism in Darwinian evolution, it is not the only one. For example, random variations in the genes transmitted from parents to offspring in a population may eventually lead to the dominance of some genes over others, a phenomenon called genetic drift. See John Beatty, "Random Drift," in Keywords in Evolutionary Biology, Evelyn Fox Keller and Elisabeth A. Lloyd, eds. (Cam-

bridge: Harvard University Press, 1992), 273-81. Similarly, if a natural catastrophe accidentally wipes out more striped animals than nonstriped animals in a population, the nonstriped survivors will dominate the surviving population, even if the gene for nonstripedness is not otherwise adaptive.

- 17. Gould, The Panda's Thumb, 83-84.
- 18. Ibid., 84.
- 19. For attempts at such an argument, see Charles J. Lumsden and Edward O. Wilson, Genes, Mind, and Culture: The Coevolutionary Process (Cambridge: Harvard University Press, 1981); William H. Durham, Coevolution: Genes, Culture, and Diversity (Stanford: Stanford University Press, 1991); Robert Boyd and Peter J. Richerson, "Why Does Culture Increase Human Adaptability?" Ethology and Sociobiology 16 (1995): 125-43.
- 20. Stephen Jay Gould and Richard C. Lewontin, "The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme," Proceedings of the Royal Society, London (1979) B. 205: 581-98.
- 21. I believe that this distinction originates with the zoologist Richard Dawkins, but I have not been able to trace the exact source.
- 22. A Lamarckian theory of evolution would also predict the existence of designoid features of organisms, because organisms might adapt favorably to their environments even if they did not do so through conscious plan or intention.
 - 23. Gould, The Panda's Thumb, 20-21, 27-29.
 - 24. Ibid., 22-26, 29.
- 25. Stephen Jay Gould, Hen's Teeth and Horse's Toes: Further Reflections in Natural History (New York: Norton, 1983), 156-57; Gould and Lewontin, "The Spandrels of San Marco," 594-97.
 - 26. Gould, The Panda's Thumb, 27-29.
- 27. Stephen Jay Gould and Elizabeth S. Verba, "Extapation: A Missing Term in the Science of Form," Paleobiology 8, no. 1 (1982): 4-15.
- 28. Gould, Hen's Teeth and Horse's Toes, 170. In his paper with Lewontin, Gould uses an example drawn from the cultural world—the spandrels in the Basilica of San Marco—to make this point. Spandrels are triangular spaces that occur when a square of four rounded archways is topped by a cathedral dome. It was customary for artists to decorate these spaces with elaborate paintings and mosaics. Nevertheless, Gould and Lewontin note, one should not infer that basilicas were specifically designed to create spandrels for artists to decorate. Instead, the custom of decorating spandrels came later; it resulted from previous decisions about the design and construction of basilicas. Gould and Lewontin, "The Spandrels of San Marco," 582-83.
- 29. Stephen Jay Gould, Ever Since Darwin: Reflections in Natural History (New York: Norton, 1977), 107-10; Gould and Verba, "Extapation," 11-12.
- 30. This is true for technology as well as for cultural software. For a description of technological bricolage see Henry Petroski, The Evolution of Useful Things (New York: Vintage, 1992).
- 31. On this point see Petroski, The Evolution of Useful Things; David Pye, The Nature and Aesthetics of Design (London: Barrie and Jenkins, 1978).

3. Memetic Evolution

- 1. Richard Dawkins, The Selfish Gene (Oxford: Oxford University Press, new ed., 1989), 192.
- 2. See Daniel C. Dennett, Consciousness Explained (Boston: Little, Brown, 1991), 200; Dawkins, The Selfish Gene, 322.
- 3. Graham Cairns-Smith, Genetic Takeover and the Mineral Origins of Life (Cambridge: Cambridge University Press, 1982); Graham Cairns-Smith, Seven Clues to the Origin of Life: A Scientific Detective Story (Cambridge: Cambridge University Press, 1985).
 - 4. Dawkins, The Selfish Gene, 192.
- 5. Robert Boyd and Peter J. Richerson, "The Evolution of Norms: An Anthropological View," Journal of Institutional and Theoretical Economics 150, no. 1 (1994): 72-87, at 74; Robert Boyd and Peter J. Richerson, Culture and the Evolutionary Process (Chicago: University of Chicago Press, 1985).
- 6. See, e.g., Daniel C. Dennett, Darwin's Dangerous Idea: Evolution and the Meanings of Life (New York: Simon and Schuster, 1995), 344; Henry C. Plotkin, Darwin Machines and the Nature of Knowledge (Cambridge: Harvard University Press, 1994), 215; Richard Brodie, Virus of the Mind: The New Science of the Meme (Seattle: Integral, 1996); Aaron Lynch, Thought Contagion: How Belief Spreads Through Society (New York: Basic, 1996). By contrast, I argue that the most basic forms of memes and meme complexes are skills.
- 7. Thus the different uses of cultural tools offered in Chapter 2 can all be redescribed as different aspects of cultural know-how. Knowing how to get about in the world, how to deal with others, and how to articulate one's values all can be and are transmitted in the form of memes.
- 8. In terms of our computer metaphor, the primacy of knowing how over knowing that tends to blur the distinction between information (data) and code (instructions). Yet this distinction is already blurred when we define memes as units of cultural transmission. If memes are to be transmitted to others, and thus become cultural, they must have some observable effects on human behavior—at the very least enough so that they are in fact transmitted. See John A. Ball, "Memes as Replicators," Ethology and Sociobiology 5 (1984): 145-61, at 154.
- 9. Dan Sperber, Explaining Culture: A Naturalistic Approach (Oxford: Blackwell 1996), 24.
 - 10. Dennett, Darwin's Dangerous Idea, 353.
 - 11. Ibid., 354.
- 12. Dan Sperber and Deirdre Wilson, Relevance: Communication and Cognition (Oxford: Blackwell, 1986).
- 13. Human beings communicate meanings to others by many devices, including signs, words, dress, and behavior. Human action is always freighted with meaning and as a result it often communicates in addition to whatever else it does. Hence we cannot restrict the notion of memetic transmission to action that is intended by the agent to communicate a message. Conversely, there is no guarantee that individuals will receive what others have deliberately sent them. Students do not always learn exactly what a teacher hopes that they will learn. They may misunderstand the teacher's lesson, learn

only parts of it, or learn nothing at all. Their ability to assimilate new information, or new ways of thinking, depends upon the cultural software they already possess. Thus it is not surprising that two persons who attend the same lesson will carry different things away from it. Their cultural software affects how each will assimilate or reject, understand or miss the point of what is being said; the cultural software of each affects how that software will change in response to what each experiences.

Our tools of understanding are also affected by interactions that are not intended to teach us anything at all. An elementary school teacher may be attempting to demonstrate how to multiply fractions, but what her pupils may be learning from her is how to dress, how to speak, and how to behave in public. If she calls on boys to solve math problems more frequently than girls, or interrupts girls more frequently than boys, they may be learning cultural lessons that the teacher may not intend for them to learn at all. The process of communicative interaction is complex and unpredictable. We therefore cannot predict how people's cultural software will be affected simply by examining what an agent intended to convey or the content of what she said. There is always a possible divergence between intended communication and effects on cultural software. Indeed, there is always the possibility that communication will have no significant effects at all.

- 14. Dennett, Consciousness Explained, 201; see Dawkins, The Selfish Gene, 195.
- 15. Dennett, Darwin's Dangerous Idea, 344.
- 16. In the key of D, the notes are f#-e'-d'-b-a-f#-a-d'-b-a.
- 17. This is because, as a practical matter, these three notes do not invoke the larger melody Tchaikovsky wrote, or the symphony as a whole, unlike, for example, the first four notes of Beethoven's Fifth Symphony. Because reproducibility and memorizability depend on environmental factors, however, there is no reason in theory why Tchaikovsky's notes could not someday be a meme.

To return to Dennett's original example, the notes D-F‡-A are the notes of a major triad, one of the building blocks of Western music. (More precisely, they are the notes of the D-major triad. This raises the interesting question of whether transpositions of melodies in different keys constitute the same meme or different memes.) These notes are reproduced continuously and reliably precisely because they are an enjoyable and satisfying combination of elements to Western ears. They are, in Dennett's words, "distinct memorable units" that music students are taught to memorize and employ in compositions. Ibid., 344. Thus they are both memes in their own right and the building blocks of other memes.

18. The U.S. Copyright Office Regulations specifically state that short phrases cannot be copyrighted. See 37 C.F.R. sec. 202.1(a) (1994) (excluding from copyright protection "words and short phrases such as names, titles, and slogans" and "familiar symbols and designs"). One reason often given for the rule is that ordinarily, short phrases do not display the creativity sufficient to justify enforcement of what is in effect a property right in their use. See, e.g., *Magic Marketing, Inc. v. Mailing Services of Pitts-burgh, Inc.*, 634 F. Supp. 769, 771 (W.D. Pa. 1986); Jessica Litman, "The Public Domain," *Emory Law Journal* 39 (1990): 965–1023, at 1013–14.

The details of intellectual property law are beyond the scope of this book. Suffice it to say that many different kinds of units, from phrases to font shapes, from techniques

to trade names, can and have been given intellectual property status under copyright, patent, or trademark laws. The layperson will likely be amazed both at the insignificance of many things that have been given intellectual property status and at the significance of many things that have been denied this status. Thus, although the statement made in the text is broadly true, it is subject to many qualifications and complications, due in part to the idiosyncracies of legislative drafting, litigation strategy, and judicial enforcement.

- 19. Ernst Mayr, The Growth of Biological Thought (Cambridge: Harvard University Press, 1982), 46-47.
- 20. Dan Sperber, "Anthropology and Psychology: Towards an Epidemiology of Representations," Man n.s. 20 (1985): 73-89, at 74.
- 21. Sanford Levinson and J. M. Balkin, "Law, Music, and Other Performing Arts," University of Pennsylvania Law Review 139 (1991): 1597-1658, at 1623.
- 22. Juan Delius, "The Nature of Culture," in The Timbergen Legacy, M. S. Dawkins, T. R. Halliday, and R. Dawkins, eds. (London: Chapman and Hall, 1991), 71-99, at 81. As Delius points out, "Culture as a persistent phenomenon is heavily dependent on long-term memories."
 - 23. Dennett, Darwin's Dangerous Idea, 348-49.
 - 24. See Dawkins, The Selfish Gene, 199.
- 25. This feature of cultural transmission underlies the deconstructive theory of the sign. The public nature of communication requires that signs be able to signify repeatedly to new users and in new contexts regardless of the intentions that originally created them. This ability of signs to be detached from the author's private intentions, and to mean something other than what the author meant, makes iterability, and hence intersubjective meaning, possible. See J. M. Balkin, "Deconstructive Practice and Legal Theory," Yale Law Journal 96 (1987): 743-86, at 779-81. As a sign is repeatedly understood, it takes on a life of its own in a relation of partial similarity and partial difference from the person who meant it. Repetition of a sign in a new context is simultaneously a relation of identity and difference; the repeated sign is syntactically identical, yet semantically different. Hence the deconstructive aphorism that "iterability alters." Jacques Derrida, "Limited Inc abc . . . ," Glyph 2 (1977): 162-254, at 200.
 - 26. Dennett, Darwin's Dangerous Idea, 353-56.
- 27. See Jon Elster, Sour Grapes: Studies in the Subversion of Rationality (Cambridge: Cambridge University Press, 1983), 152–53.
 - 28. See the discussion of cognitive dissonance theory in Chapter 8.
- 29. Charles Darwin, On the Origin of Species by Means of Natural Selection, in The Portable Darwin, Duncan M. Porter and Peter W. Graham, eds. (New York: Penguin, 1993), 185; Niles Eldredge, Reinventing Darwin: The Great Debate at the High Table of Evolutionary Theory (New York: Wiley, 1995), 50.
- 30. James Burke, Connections (Boston: Little, Brown, 1978), 108-13. This book, based on the television series of the 1970s, contains many wonderful examples of technological borrowing.
- 31. See Stephen Jay Gould, "The Panda's Thumb of Technology," in Bully for Brontosaurus (New York: Norton, 1991), 59-75, at 65. The ability of memes to combine

in human minds means that cladistics—the study of lineages—is extremely difficult in the case of memes. Yet it is an important element in the study of biological evolution. Eldredge, *Reinventing Darwin*, 53–55. One might think of intellectual history as a sort of cladistics of memes. Intellectual historians often try to study ideas as they change through history, but the theory of memes suggests why this enterprise presents so many complications.

32. Dan Sperber calls this an "attraction model" of cultural evolution because the transformation of cultural software (or cultural representations, as he calls them) tends to converge toward the most popular versions, or "attractors." The term *attractor* is borrowed from chaos theory. An attractor attracts nothing; it is simply the standard set of features toward which successive transformations tend, on the average, to move. Once near an attractor, subsequent transformations tend to stay in the general vicinity. Sperber notes that the reasons why transformations converge on attractor points may depend on universal features of human psychology or the vicissitudes of the local cultural environment. Changes in the cultural environment may shift cultural attractor points and lead to large-scale shifts in belief and practice. Sperber, *Explaining Culture*, 105–18.

Sperber contrasts his attraction model to the evolutionary theories of Dawkins and Dennett, which focus on the survival of memes in individual minds rather than on their successive transformation. He also rejects the use of the word *meme* because he assumes that memetic evolution necessarily presupposes virtually exact copying of cultural information, employing human beings as mere "agents of replication... with little or no individual contribution to the process" (105–6). This strikes me as a bit of a caricature. A Darwinian theory of cultural evolution is in no sense committed to this position. Sperber is engaged in a play on words, identifying the word *meme* with its French cognate (*même*, meaning same or identical) rather than focusing on its connection with memory. Memories of events surely change as they are transmitted from person to person; memory is rarely, if ever, *la même chose*.

There is no reason why the use of the term *meme* has to be tied to the fallacious assumption that cultural transmission is a matter of perfect copying. Theories of cultural evolution should be based on the recognition that although sometimes replication of cultural information is fairly exact, more often it is not. Scribes may carefully copy manuscripts, but musical performers improvise. Xerox machines duplicate, but cooks change proportions and add new ingredients. Indeed, symbolic forms that exist outside human minds are much more likely to be exact copies of each other than the cultural software in human minds. That is because human technology can create exact copies, while the processes of human memorization and understanding rarely do.

Whatever we call the units of cultural transmission, whether memes, representations, or something else, a theory of cultural evolution must reckon with both differential rates of attractiveness to other minds and distinctive forms of transformation by the minds who possess them. This is Sperber's deeper point. Cultural evolution must be shaped not only by those factors that ensure the survival of descendants but by those factors that ensure that the descendants remain roughly similar to each other.

33. See Dennett, *Consciousness Explained*, 204: Donald Campbell, "Comments on the Sociobiology of Ethics and Moralizing," *Behavioral Science* 24 (1979): 37–45.

- 34. Dennett, Darwin's Dangerous Idea, 348. As Dennett points out, "Plato's ideas survive not because of the survival of individual papyrus manuscripts, but because they were continuously copied."
- 35. See Donald R. Griffin, Animal Minds (Chicago: University of Chicago Press, 1992); John Tyler Bonner, The Evolution of Culture in Animals (Princeton: Princeton University Press, 1980); Merlin Donald, Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition (Cambridge: Harvard University Press, 1991).
- 36. Thus birdsongs are a kind of meme that can survive in the environment that bird's minds and bodies provide. See Dawkins, The Selfish Gene, 189-90.
 - 37. Dennett, Consciousness Explained, 202.
 - 38. Ibid., 202, 206.
 - 39. See Dennett, Consciousness Explained, 218.
 - 40. Ibid., 220.
 - 41. Ibid.
 - 42. Dennett, Darwin's Dangerous Idea, 350.
- 43. Richard Nisbett and Lee Ross, Human Inference: Strategies and Shortcomings of Social Judgment (Englewood Cliffs, N.J.: Prentice-Hall, 1980), 169-88.
 - 44. See the discussion in Chapter 8.
 - 45. Dennett, Darwin's Dangerous Idea, 351.
 - 46. Ibid.
- 47. Sperber, "Anthropology and Psychology," 74. Cavalli-Sforza and Feldman have explicitly attempted to model cultural transmission on the transmission of disease. L. L. Cavalli-Sforza and M. W. Feldman, "Models for Cultural Inheritance," 1, "Group Mean and Within Group Variation," Theoretical Population Biology 4 (1973): 42-45; L. L. Cavalli-Sforza and M. W. Feldman, Cultural Transmission and Evolution: A Quantitative Approach (Princeton: Princeton University Press, 1981); L. L. Cavalli-Sforza and M. W. Feldman, "Cultural Versus Genetic Adaptation," Proceedings of the National Academy of Sciences, USA 80 (1983): 4993-96. Their models also use the idea of genetic drift as an evolutionary mechanism.

Similar models have been proposed for the spread of technological innovation. Everett M. Rogers, Diffusion of Innovations (New York: Free Press, 3d ed., 1983). The cumulative adoption of an innovation usually seems to follow an S-shaped curve which resembles the spread of communicable diseases.

- 48. Sperber, "Anthropology and Psychology," 74. Brodie, Virus of the Mind, and Lynch, Thought Contagion, are two recent accounts of memetics that are organized around the communicable-disease analogy. Each offers abundant historical and cultural examples.
 - 49. Dennett, Darwin's Dangerous Idea, 352.
 - 50. Sperber, "Anthropology and Psychology," 74.
 - 51. Delius, "The Nature of Culture," 84.
 - 52. Ibid., 84-85.
 - 53. Ibid., 86-87.
 - 54. Ibid., 87.

- 55. Charles J. Lumsden and Edward O. Wilson, Genes, Mind, and Culture: The Coevolutionary Process (Cambridge: Harvard University Press, 1981).
- 56. Dennett, Consciousness Explained, 203; Richard Dawkins, The Extended Phenotype (San Francisco: Freeman, 1982), 110-11.
- 57. Conspiracy theories and beliefs that are linked to prohibitions against exposing one's self to contrary beliefs are also good examples of self-reinforcing beliefs. See Dennett, Consciousness Explained, 206; Dawkins, The Selfish Gene, 198-99.
- 58. Stephen Jay Gould, Hen's Teeth and Horse's Toes (New York: Norton, 1983), 174-75.
 - 59. Ibid., 167.
 - 60. Ibid., 175.
 - 61. Ibid., 173.
 - 62. Lumsden and Wilson, Genes, Mind, and Culture, 13.
- 63. The fact that so many members of the Catholic clergy have been able to suppress their reproductive urges over so many years is perhaps the best evidence of the power of memes over genes. Even though Catholic clergy occasionally (and predictably) have gone astray, the degree of celibacy that they have been able to practice over the centuries is certainly remarkable.
- 64. An example is the development of strains of bacteria that are resistant to antibiotics—although this is a response not to human biological evolution but to changes in human technology.
- 65. George C. Williams and Randolph M. Neese, "The Dawn of Darwinian Medicine, Quarterly Review of Biology 66, no. 1 (March 1991): 1-22, at 7.
 - 66. Delius, "The Nature of Culture," 86, 91.
 - 67. Ibid., 85.
 - 68. Ibid., 89.
 - 69. Ibid., 93.
 - 70. Ibid., 89-90.
- 71. See Alexander Rosenberg, "Altruism: Theoretical Contexts," in Keywords in Evolutionary Biology, Evelyn Fox Keller and Elisabeth A. Lloyd, eds. (Cambridge: Harvard University Press, 1992), 19-28.
 - 72. Delius, "The Nature of Culture," 93-94.
 - 73. Ibid., 94.
- 74. Because memetic kinship can differ from genetic kinship, sometimes cultural and biological altruism will be reinforcing and sometimes they will be at odds. The bond between parents and children is no doubt strengthened by the cultural transmission that usually accompanies parenting. On the other hand, there are many stories of families split asunder by civil wars and religious disputes.
 - 75. Ball, "Memes as Replicators," 156.
 - 76. Sperber, "Anthropology and Psychology," 86.
- 77. For a helpful discussion see Roger C. Schank, The Connoisseur's Guide to the Mind: How We Think, How We Learn, and What It Means to Be Intelligent (New York: Summit, 1991), 35-41.

- 78. See, e.g., Ken Binmore, Game Theory and the Social Contract: Playing Fair, vol. 1 (Cambridge: MIT Press, 1994); Robert Axelrod, The Evolution of Cooperation (New York: Basic, 1984); Edna Ullmann-Margalit, The Emergence of Norms (Oxford: Clarendon, 1977).
- 79. Stephen Jay Gould and Richard C. Lewontin, "The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme," *Proceedings of the Royal Society, London* (1979) B.205: 581–98.
 - 80. Eldredge, Reinventing Darwin, p. 46.
- 81. For examples of this point in the evolution of accident law, see J. M. Balkin: "Too Good to Be True: The Positive Economic Theory of Law," *Columbia Law Review* 87 (1987): 1447–89.
 - 82. Gould and Lewontin, "The Spandrels of San Marco," 582-83.

4. The Spread of Cultural Software

- 1. See John A. Ball, "Memes as Replicators," *Ethology and Sociobiology* 5 (1984): 145-61.
- 2. Daniel C. Dennett, *Darwin's Dangerous Idea: Evolution and the Meanings of Life* (New York: Simon and Schuster, 1995), 349; see also Richard Dawkins, "Viruses of the Mind," in *Dennett and His Critics*, Bo Dahlbom, ed. (Oxford: Blackwell, 1993), 13–27.
- 3. Richard Dawkins, *The Selfish Gene* (Oxford: Oxford University Press, new ed., 1989), 212.
 - 4. Dennett, Darwin's Dangerous Idea, 349.
 - 5. Ibid.
- 6. Dan Sperber, "The Epidemiology of Beliefs," in *The Social Psychological Study of Widespread Beliefs*, Colin Fraser and George Gaskell, eds. (Oxford: Clarendon, 1990), 25–44.
 - 7. Dennett, Darwin's Dangerous Idea, 349.
- 8. See Dan Sperber, "Anthropology and Psychology: Towards an Epidemiology of Representations," *Man* n.s. 20 (1985): 73–89, at 82.
 - 9. Ibid., 80-83.
- 10. Eric A. Havelock, *Preface to Plato* (Cambridge: Harvard University Press, 1963); Albert B. Lord, *The Singer of Tales* (Cambridge: Harvard University Press, 1960).
- 11. Sperber, "Anthropology and Psychology," 86. Sperber calls this the "Law of the Epidemiology of Representations" for oral cultures.
- 12. Neil Postman, Amusing Ourselves to Death: Public Discourse in the Age of Show Business (New York: Penguin, 1985); Marshall McLuhan, Understanding Media: The Extensions of Man (New York: McGraw-Hill, 1964); Marshall McLuhan, The Gutenberg Galaxy: The Making of Typographic Man (Toronto: University of Toronto Press, 1962).
- 13. Walter J. Ong, *Orality and Literacy: The Technologizing of the Word* (London: Methuen, 1982).
- 14. Postman, Amusing Ourselves to Death; McLuhan, The Gutenberg Galaxy; Ong, Orality and Literacy, 135–38.

- 15. To vary McLuhan's famous aphorism, the medium is the meme as well as the message.
- 16. Postman, Amusing Ourselves to Death; Ronald Collins and David Skover, The Death of Discourse (New York: HarperCollins, 1996), J. M. Balkin, "What Is a Postmodern Constitutionalism?" Michigan Law Review 92 (1992): 1966-90.
- 17. Over time the average length of uninterrupted statements of presidential candidates in the United States has been shrinking, and so has the length of campaign advertisements. Cass Sunstein, Democracy and the Problem of Free Speech (New York: Free Press, 1993), 61; Kathleen Hall Jamieson, Dirty Politics (New York: Oxford University Press, 1992), 205–8.
- 18. For a particularly pessimistic version of this thesis, see Postman, Amusing Ourselves to Death.
 - 19. Sperber, "Anthropology and Psychology," 80-81.
- 20. In historical linguistics, for example, Grimm's Law predicts the direction in which pronunciation of consonants will mutate over time. See Theodora Binyon, Historical Linguistics (Cambridge: Cambridge University Press, 1977), 83-85.
 - 21. Sperber, "Anthropology and Psychology," 75.
- 22. Sperber, "Epidemiology of Beliefs," 29-30. This is consistent with experiments which show that narratives tend to be pared down and simplified as they are transmitted from person to person.
 - 23. Ibid.
- 24. "Then Midianite traders passed by; and they drew Joseph up and lifted him out of the pit, and sold him to the Ishmaelites for twenty shekels of silver; and they took Joseph to Egypt." Gen. 37:28, Revised Standard Version. Interestingly, Joseph himself later accuses his brothers of having sold him into slavery. See Gen. 45:4 ("I am your brother, Joseph, whom you sold into Egypt"). The potential ambiguity in the biblical description of the sale stems from the use of the Hebrew word vayimk'ru ("and they sold") appearing after the reference to the Midianites; however, most English versions of the Bible—including the King James, the New American Standard, the Jewish Publication Society, and the New English Bible—translate the passage similarly to the Revised Standard. It is possible that the biblical text is a conflation of two versions of the story—yet another example of the phenomenon of mutations in storytelling.

Joseph's sale into slavery has posed an interpretive problem for biblical commentators. Rashi's commentary on the Torah argues that Jacob's sons sold Joseph to the Ishmaelites, who sold him to the Midianites, who sold him to the Egyptians, who brought Joseph into Egypt. Thus the phrasing of the Hebrew in Gen. 37:28 is meant to suggest not that Joseph was not sold by his brothers, but that Joseph was sold many times before he arrived in Egypt. The Pentateuch and Rashi's Commentary, Rabbi Abraham Ben Isaiah and Rabbi Benjamin Sharfman, trans., vol 1 (Brooklyn: S.S. and R., 1949), 379. Interestingly, the New International Version simply avoids the textual conflict altogether by reading the expression "and they sold" to refer to the brothers. Gen. 37:28 New International Version.

25. For a collection of such transformed quotations, see Paul F. Boller Jr. and John

George, They Never Said It: A Book of Fake Quotes, Misquotes, and Misleading Attributions (New York: Barnes and Noble, 1989).

- 26. Claude Lévi-Strauss, *The Raw and the Cooked: Introduction to a Science of Mythology*, vol. 1, John Weightman and Doreen Weightman, trans. (New York: Octagon, 1970).
 - 27. Ball, "Memes as Replicators," 155.
 - 28. Sperber, "Epidemiology of Beliefs," 33-34.
 - 29. Ball, "Memes as Replicators," 155.
 - 30. Sperber, "Anthropology and Psychology," 84.
- 31. Human knowledge often uses coherence as an organizational principle; beliefs are often rejected if they do not square fully with beliefs already held. However, because human beings can hold beliefs that they do not completely understand, they may avoid rejecting these beliefs until more information arrives that might make their beliefs coherent or produce a better understanding of them. They may hold some beliefs because, only half-understanding them, or lacking knowledge as to whether they are true, they simply take the beliefs on authority. See Sperber, "Epidemiology of Beliefs," 33–34. Thus a person may take on authority both the half-understood belief that space curves near a heavy mass and the half-understood belief that a communion wafer is transubstantiated into the body of Jesus, even though the sources of authority and the institutionally recognized justifications for the two half-understood beliefs differ. As a result, people may be able to hold a number of beliefs that are seemingly in tension with each other, because the grounds of belief for each are of a different status.
- 32. Sperber, "Epidemiology of Beliefs," 36–37; Sperber, "Anthropology and Psychology," 84–85.
- 33. Juan Delius, "The Nature of Culture," in *The Timbergen Legacy*, M. S. Dawkins, T. R. Halliday, and R. Dawkins, eds. (London: Chapman and Hall, 1991), 71–99, at 95.
- 34. Helena Cronin, "Sexual Selection: Historical Perspectives," in Evelyn Fox Keller and Elisabeth A. Lloyd, eds., Keywords in Evolutionary Biology (Cambridge: Harvard University Press, 1992), 286–93. Indeed, the preferred trait may actually be a handicap; the standard example is the male peacock's tail, which is a greater burden the longer and more gaudy it becomes. Amotz Zahavi has argued that these self-imposed handicaps may actually serve as a positive signal for mate selection: If a male peacock can successfully drag around a ridiculously long tail, he must be very fit indeed. Thus females will gravitate to the most handicapped males as long as they are able to survive and mate. Amotz Zahavi, "The Theory of Signal Selection and Some of Its Implications," in V. P. Delfino, ed., International Symposium on Biological Evolution, Bari, 9–14 April 1985 (Bari, Italy: Adriatici Editrici), 305–27; Amotz Zahavi, "Mate Selection: A Selection for a Handicap," Journal of Theoretical Biology 53 (1975): 205–14.
- 35. Ball, "Memes as Replicators," 151. See also Robert Boyd and Peter J. Richerson, Culture and the Evolutionary Process (Chicago: University of Chicago Press, 1985), 259–79.
 - 36. Delius, "The Nature of Culture, 95-96.
 - 37. Dennett, Darwin's Dangerous Idea, 352.
 - 38. Sperber, "Epidemiology of Beliefs," 37.
- 39. Robert Boyd and Peter J. Richerson, "The Evolution of Norms: An Anthropological View," *Journal of Institutional and Theoretical Economics* 150, no. 1 (1994): 72-

- 87; Peter J. Richerson and Robert Boyd, "Darwinian Models of Culture: Toward Replacing the Nature/Nurture Dichotomy," World Futures 34 (1991): 43-57, at 50-52.
 - 40. Sperber, "Epidemiology of Beliefs," 41.
 - 41. 1163 U.S. 537 (1896).
- 42. Martin Luther King, Jr., A Testament of Hope: The Essential Writings of Martin Luther King, 7r., James Melvin Washington, ed. (San Francisco: Harper and Row, 1986), 219.
- 43. See J. M. Balkin, "Some Realism About Pluralism: Legal Realist Approaches to the First Amendment," Duke Law Journal 1990: 375-430.
- 44. Ibid.; J. M. Balkin, "Ideological Drift and the Struggle over Meaning," Connecticut Law Review 25 (1993): 875-91; J. M. Balkin, "Ideological Drift," in Roberta Kevelson, ed., Action and Agency: Fourth Round Table on Law and Semiotics (New York: Peter Lang, 1991).
- 45. In this way we can offer an evolutionary account of Bourdieu's "economy of logic" discussed in Chapter 2.
- 46. Robert C. Ellickson, "Property in Land," Yale Law Journal 102 (1993): 1315-1400.
 - 47. Dennett, Darwin's Dangerous Idea, 486.
- 48. Thus it was not accidental that Louis Althusser identified them as examples of "ideological state apparatuses." Louis Althusser, "Ideology and Ideological State Apparatuses (Notes Towards an Investigation)," in Lenin and Philosophy and Other Essays (New York: Monthly Review Press, 1971), 127-86.
- 49. See Ernst Mayr, Toward a New Philosophy of Biology: Observations of an Evolutionist (Cambridge: Harvard University Press, 1988), 318-19; Ernst Mayr, The Growth of Biological Thought (Cambridge: Harvard University Press, 1982), 270-75.
- 50. Roger C. Schank, The Connoisseur's Guide to the Mind: How We Think, How We Learn, and What It Means to Be Intelligent (New York: Summit, 1991), 41. Schank argues that the most distinctive cooking styles are often those of communities where isolation has led to rigidification of expectations about how food should be prepared.
- 51. Emile Durkheim, The Division of Labor in Society (New York: Free Press, 1964), 167-73.
- 52. See Sanford Levinson and J. M. Balkin, "Law, Music, and Other Performing Arts," University of Pennsylvania Law Review 139 (1991): 1597-1658, for an account of the authentic performance movement along these lines.

5. Conceptions of Ideology

- 1. Jon Elster, Making Sense of Marx (Cambridge: Cambridge University Press, 1985), 462-64.
- 2. John Thompson, Ideology and Modern Culture (Stanford: Stanford University Press, 1991), 59 ("By 'symbolic forms' I understand a broad range of actions and utterances, images and texts, which are produced by subjects and recognized by them and others as meaningful constructs"); Clifford Geertz, The Interpretation of Cultures (New

York: Basic, 1973), 212–15. The locus classicus of the term is Ernst Cassirer, *The Philosophy of Symbolic Forms*, vol. 1, *Language* (New Haven: Yale University Press, 1955).

3. In contrast, Thompson does not seem concerned with a distinction between mental processes and symbolic forms existing in the social world. His approach is sociological rather than philosophical, and hence these questions do not concern him. On the other hand, in his essay "Ideology as a Cultural System," Geertz seems specifically interested in externalizing the study of ideology from internal mental operations to symbols. Geertz, *The Interpretation of Cultures*, 215.

It is interesting to note that Cassirer's original use of this concept was strongly Kantian in spirit. Cassirer argued that symbolic forms in science, language, myth, art, and religion constructed the world for us and enabled us to understand it; at the same time he viewed these forms as functions of mind that allowed people to conceive both the world and themselves and created a bridge between the two. Cassirer, *The Philosophy of Symbolic Forms*, 1: 91.

- 4. See, e.g., Catharine MacKinnon, Feminism Unmodified: Discourses on Life and Law (Cambridge: Harvard University Press, 1987), 36, 42–43, 71.
- 5. Michéle Barrett, The Politics of Truth: From Marx to Foucault (Cambridge: Polity, 1991), 4.
- 6. See Jon Elster, Sour Grapes: Studies in the Subversion of Rationality (Cambridge: Cambridge University Press, 1983), 163.
 - 7. Elster, Making Sense of Marx, 4, 27-29.
 - 8. Ibid., 18-27, 465-93.
 - 9. Thompson, Ideology and Modern Culture, 56.
- 10. See Ben-Ami Shillony, *The Jews and The Japanese: The Successful Outsiders* (Rutland, Vt.: Tuttle, 1991), 216-22.
- 11. Sheila K. Johnson, "Japanese and Jews: Intersection of Myths," Los Angeles Times, November 30, 1992. A good example of this mixture of admiration and negative stereotyping is a 1972 book by Fujita Den, the president of McDonald's of Japan. Entitled Jewish Trade Practices, it advises the Japanese to be more shrewd and unscrupulous like Jews in order to achieve business success. Shillony, The Jews and The Japanese, 217.
- 12. Leslie Helm, "Japan Newspaper Ad Revives Fears of Anti-Semitism," Los Angeles Times, July 29, 1993.
 - 13. Elster, Sour Grapes, 157.
 - 14. Thompson, Ideology and Modern Culture, 56, 59.
 - 15. Ibid., 68, 73.
- 16. For example, the psychological mechanisms that reduce cognitive dissonance may operate to produce self-serving justifications by members of subordinate groups toward each other and forms of wishful thinking that work to the advantage of members of relatively dominant groups, as discussed in Chapter 8.
- 17. Catharine MacKinnon, Toward a Feminist Theory of the State (Cambridge: Harvard University Press, 1989), 116.
- 18. Thompson's model of ideology does not specifically consider the competing ideologies of subordinated groups. He is concerned only with the counterideology of

subordinate groups that he calls "incipient forms of the critique of ideology." See Thompson, Ideology and Modern Culture, 68.

19. See, e.g., Elizabeth V. Spellman, Inessential Woman: Problems of Exclusion in Feminist Thought (London: Women's Press, 1988); Martha R. Mahoney, "Whiteness and Women, in Practice and Theory: A Reply to Catharine MacKinnon," Yale Journal of Law and Feminism 5 (1993): 217-51.

Ambivalence and Self-Reference

- 1. See Michéle Barrett, The Politics of Truth: From Marx to Foucault (Cambridge: Polity, 1991), 19; Jorge Larrain, The Concept of Ideology (London: Hutchison, 1979).
 - 2. See, e.g., Barrett, The Politics of Truth, 18-26.
- 3. Lukacs also sometimes speaks of bourgeois consciousness as "false." See, e.g., Georg Lukacs, History and Class Consciousness: Studies in Marxist Dialectics (Cambridge: MIT Press, 1971), 53-54. Jorge Larrain has insisted that despite these remarks, Lukacs's basic conception of ideology is neutral because it does "not pass judgment on the validity or adequacy of ideas." Jorge Larrain, Marxism and Ideology (London: Macmillan, 1983), 73, 239 n. 73.
- 4. Raymond Geuss, The Idea of a Critical Theory: Habermas and the Frankfurt School (Cambridge: Cambridge University Press, 1981), 22-26.
- 5. V. I. Lenin, "What Is to Be Done?" in The Lenin Anthology, Robert C. Tucker, ed. (New York: Norton, 1975), 50.
- 6. John Thompson, Ideology and Modern Culture (Stanford: Stanford University Press, 1991), 95; Clifford Geertz, "Ideology as a Cultural System," in The Interpretation of Cultures (New York: Basic, 1973), 193-233.
- 7. Karl Mannheim, Ideology and Utopia: An Introduction to the Sociology of Knowledge (New York: Harcourt Brace Jovanovitch, 1936), 120-24, 149-51.
 - 8. Ibid., 88-89, 93-94.
- 9. The expression "Mannheim's paradox" comes from Geertz, The Interpretation of Cultures, 194. I borrow this particular formulation of the paradox from Paul Ricoeur, Lectures on Ideology and Utopia, George H. Taylor, ed. (New York: Columbia University Press, 1986), 157.
 - 10. Terry Eagleton, Ideology: An Introduction (London: Verso, 1991), 2.
- 11. Mannheim, Ideology and Utopia, 76-77. The recognition of this symmetry produces what Mannheim calls the general conception of ideology; Mannheim argues that this development transforms the simple theory of ideology into the sociology of knowledge (77–78).
 - 12. See, e.g., his discussion of fascism, ibid., 134-46.
- 13. This is in accord with Gadamer's argument about the necessity of prejudgments and foreunderstandings as preconditions for understanding an Other, whether this Other is a text or a person. Hans-Georg Gadamer, Truth and Method (New York: Crossroad, 1975), 245–67.
- 14. The obligations imposed by the ambivalent conception are similar to those imposed by the hermeneutic circle as recast in Gadamer's ontological hermeneutics. Gad-

amer argues that we must assume that a text has some truth to convey to us; otherwise we cannot be certain whether our conclusions that the text is false or incoherent are due to the text or to our misunderstanding of it. See Gadamer, *Truth and Method*, 261–63. For a helpful discussion, see Georgia Warnke, *Gadamer: Hermeneutics, Tradition, and Reason* (Stanford: Stanford University Press, 1987), 86–91. Like Gadamer's hermeneutics, the ambivalent conception of ideology recognizes that human beings are fallible and finite creatures produced by circumstance and existing in a historical moment. This recognition produces the duty of understanding imposed by the hermeneutic circle.

Although there are important similarities, it is important not to confuse this argument for hermeneutic charity with Donald Davidson's arguments for radical interpretation according to the Principle of Charity. To begin with, Davidson is attempting a very different sort of project. He is trying to construct a theory of meaning: he wants to give an account of what goes on when a person means something in a natural language like English. Davidson argues that the idea of meaning is inextricably tied to concepts of belief, rationality, and truth. Because Davidson argues that truth is a primitive concept, he explains meaning, belief, and rationality in terms of it.

Davidson argues that when we try to understand another person, our sense of what she believes, the truth of what she is saying, and the meaning of what she is saying are interdependent. If we vary the meaning of the words, our ascriptions of belief to the speaker and the truth of those beliefs will also vary. Hence his Principle of Charity holds the truth of other people's beliefs constant and then interprets their meaning in light of this assumption. "This is accomplished by assigning truth conditions to alien sentences that make native speakers right when plausibly possible, according, of course, to our own view of what is right." Donald Davidson, *Inquiries into Truth and Interpretation* (Oxford: Clarendon, 1984), 137.

The task of ideological analysis is quite different. In ideological analysis, we assume (1) that others have beliefs quite different from ours, (2) that much of what we regard as obvious they do not, and (3) that part of our task is to learn something new from them in the task of critiquing them. The goal of Davidson's radical interpretation is not to have his ethics or political morality enlightened by an encounter with the natives; it is rather to provide a translation manual for their language. Radical interpretation does not seek substantive knowledge about what is good or true but rather semantic knowledge about what the terms of another person's language mean.

Although Davidsonian interpretation seems charitable in that it tries to make the statements of other persons true, it is really a charity that begins at home. It assumes that our own beliefs are true and then tries to vary the meanings of what others are saying so that their statements conform to our beliefs. The hermeneutic charity required by ideological analysis assumes that the truth lies somewhere between ourselves and the analysand and that it is our job to discover it. Davidsonian charity does not put our own beliefs into question to interpret the beliefs of others; hermeneutic charity requires that we do so.

Finally, Davidson's theory of meaning and his paradigmatic examples are primarily directed to questions of fact rather than questions of value. If a native reports that the statuette on my desk is "ugly" and I believe it is beautiful, Davidson does not apply the

Principle of Charity to conclude that the word ugly means "beautiful" to the native. Rather, Davidson suggests that we would accommodate this evidence in other ways; we would call this a "difference of opinion" (197). Such "differences of opinion" are often conflicts of values and value judgments. They are the primary concern of the hermeneutic charity involved in ideological analysis.

- 15. Hence an important difference between a critical approach and Gadamer's hermeneutics is that we do not engage in this approach with the goal of reaching an agreement with the analysand. Rather, we are interested in discovering both what we can learn from the analysand and what we ultimately cannot agree with because of the ideological effects we perceive in the analysand's thought.
- 16. The analogous point in the hermeneutical tradition is Gadamer's claim that understanding requires hermeneutic openness to the truth contained within a text. Gadamer, Truth and Method, 262; Warnke, Gadamer: Hermeneutics, Tradition, and Reason, 89.
- 17. These phenomena are discussed more fully, and in the context of legal judgments, in J. M. Balkin, "Understanding Legal Understanding: The Legal Subject and the Problem of Legal Coherence," Yale Law Journal 103 (1993): 105-76.
- 18. Hermeneutic co-optation is an obvious danger in Gadamer's theory of understanding because he insists that understanding seeks not only openness to but also agreement with the Other; see, e.g., Warnke, Gadamer: Hermeneutics, Tradition, and Reason, 90-91. It would be more correct to say that Gadamer's account of understanding permits the phenomena of hermeneutic conformation and co-optation as well as more benign forms of understanding. In short, Gadamer gives us an account of understanding that, while designed to show how understanding is possible, also shows how various ideological effects in our understanding can occur. For further discussion see Balkin, "Understanding Legal Understanding," 159-66.
 - 19. Mannheim, Ideology and Utopia, 77.
 - 20. Thompson, Ideology and Modern Culture, 49.
 - 21. Mannheim, Ideology and Utopia, 153-64.
- 22. Stanley Fish, Doing What Comes Naturally: Change, Rhetoric, and the Practice of Theory in Literary and Legal Studies (Durham, N.C.: Duke University Press, 1989), 436-67.
- 23. For the most succinct statement of this ubiquitous trope in Fish's work, see Stanley Fish, There's No Such Thing as Free Speech (and It's a Good Thing Too) (Oxford: Oxford University Press, 1994), 295-96.
- 24. Ernest Gellner, Reason and Culture (Cambridge: Cambridge University Press, 1992), 132.

7. Transcendence

1. The most obvious candidate for a transcendent value other than truth and justice would be beauty, although it is unclear to what extent aesthetic order and normative order are fully separate in many different cultures, including our own.

My colleague Owen Fiss has suggested to me that human solidarity is also a transcendent value. Solidarity, however, is only a special case of a more fundamental value,

which is love. The ancient Greeks divided the concept of love into affection and sexual attraction (eros), friendship (philia), and concern for the well-being of others (agapē).

There is much to recommend the notion that love is a transcendent value. Socrates' famous speech in Plato's *Symposium* views love as an unfulfilled (and unfulfillable) longing. I would argue that in this speech Plato offers us an *erotics* of human values. This erotic theory models human values on the example of love; it argues that human values are an inchoate and always unfulfilled longing and searching for the Good. This erotic conception of human values is as profound as anything Plato offers us in his middle dialogues.

- 2. Here again we should note the potential distinction between subjects of justice, who can be treated unjustly, and agents of justice, who can act unjustly.
- 3. This is one reason, I think, why theorists like Bruce Ackerman and Jürgen Habermas have turned to idealized forms of dialogue as means of explicating concepts of justice and truth. See Bruce Ackerman, Social Justice in the Liberal State (New Haven: Yale University Press, 1980); Bruce Ackerman, "Why Dialogue?" Journal of Philosophy 86 (1989): 5–22; Jürgen Habermas, Knowledge and Human Interests, Jeremy J. Shapiro, trans. (Boston: Beacon, 1971); Jürgen Habermas, Legitimation Crisis, Thomas McCarthy, trans. (Boston: Beacon, 1975); Jürgen Habermas, The Theory of Communicative Action, Thomas A. McCarthy, trans. (Boston: Beacon, vol. 1, 1984; vol. 2, 1987). Like other philosophical theories, dialogic theories of justice and truth are articulations of our transcendent ideals. As articulations, they presuppose the existence of transcendent ideals rather than produce them.
- 4. Habermas has argued that certain ideal criteria are presupposed in communicative encounters; he has tried to capture them in his notion of an "ideal speech situation." Jürgen Habermas, "Discourse Ethics: Notes on Philosophical Justification," in *The Communicative Ethics Controversy*, Seyla Benhabib and Fred Dallmayr, eds. (Cambridge: MIT Press, 1990), 60–110, at 85; Jürgen Habermas, "Warheitstheorien," in H. Fahrenbach, ed., *Festschrift für W. Schultz* (Pfullingen: Neske, 1973), 211–65; Habermas, *Legitimation Crisis*, 110. Thus Habermas argues that "participants in communication cannot avoid the presupposition that the structure of their communication . . . rules out all external or internal coercion other than the force of the better argument, and thereby also neutralizes all motives other than that of the cooperative search for truth." Habermas, "Discourse Ethics," 86.

My argument differs from Habermas's in two important respects. First, Habermas relies on procedural and substantive criteria of an ideal speech situation instead of transcendent ideals of truth and justice. Indeed, he tries to derive ideals of factual and moral truth from the results of an ideal rational consensus. Later in this chapter I shall argue that a theory of ideal consensus presupposes these transcendent ideals and that an ideal speech situation is at best a heuristic for articulating them.

Second, I do not believe that when people engage in discourse they must presume that their discourse either does or can approximate the criteria of an ideal speech situation. I seriously doubt whether the notion of an ideal speech situation involving finite human beings with limited perspectives and historically generated cultural software is a coherent one. If the idea is incoherent, there is no reason to think that it is presumed in people's speech acts.

- 5. Sometimes we and the analysand will agree totally about what is right in a particular situation. Even so, we must still acknowledge that our views of what is just are revisable, incomplete, and imperfect. From a larger perspective what we think to be unambiguously just may be much more complicated and problematic. To acknowledge this we must still postulate a regulative ideal of justice against which our current judgments might be found wanting.
- 6. Many philosophers have advanced various versions of moral relativism. See Gilbert Harman and Judith Jarvis Thompson, Moral Relativism and Moral Objectivity (Oxford: Blackwell, 1996); David B. Wong, Moral Relativity (Berkeley: University of California Press, 1984). But these theories often make some accommodation for principles of tolerance and for possibilities of moral dialogue between peoples of different cultures. For example, Wong argues that certain moral principles of tolerance apply to all agents even if they are not "universally justifiable to all agents" (189). So his theory is not strongly relativist in the sense that I discuss in the text.

Harman comes closer to that position. He argues that people can evaluate the actions of a person either relative to their own values or relative to the values of the other person. Similarly, we can either offer reasons for action that make sense from our perspective or offer reasons that would carry weight with the other person. But there is no transcultural notion of morality. It is true that many people believe in tolerating the views of others, but if a principle of tolerance is widespread, it is because from different perspectives many people have good reasons to abide by it; it is not because a principle of tolerance applies to all people generally.

Harman recognizes only one way of criticizing the views of others if what they did was right from their perspective: Harman argues that although we cannot say that it was wrong of a person to do an act that is consistent with that person's values but not our own, we can properly say that it was wrong that the person did the act. We can say that what a person did was wrong in the same way that we can say that it was bad that a tiger mauled children, or that it was bad that an enemy took steps that worked against our interests. Thus, although it makes no sense to say that it was wrong of Hitler to exterminate Jews (assuming that Hilter had good reasons from within his own value system), we can say that it was wrong that Hitler exterminated Jews (49, 59-61). It is by no means clear how much work Harman thinks this distinction can do in dealing with problems of justice between cultures.

- 7. Jean-François Lyotard and Jean-Loup Thébaud, Just Gaming (Minneapolis: University of Minnesota Press, 1985), 100.
 - 8. Ibid.
- 9. This argument is taken from J. M. Balkin, "Transcendental Deconstruction, Transcendent Justice," Michigan Law Review 92 (1994): 1131-86, at 1175.
- 10. The most famous ideal process theory is John Rawls's theory of the original position. John Rawls, A Theory of Justice (Cambridge: Harvard University Press, 1971).
 - 11. The best examples are Jürgen Habermas's and Bruce Ackerman's work.
- 12. Charles Sanders Peirce's view of truth as the eventual consensus of a community of investigators can also be understood as an ideal process theory because the consensus is never the actual consensus of any given time but is always deferred. See Charles

Sanders Peirce, "How to Make Our Ideas Clear," in *Philosophical Writings of Peirce*, Justus Buchler, ed. (New York: Dover, 1955), 23-41, at 38.

13. John Rawls's theory of the original position is probably not an ideal dialogic theory in the sense I have described. Rawls imagines his participants in mutual discussion, but the principles of justice that emerge are not contingent upon the actual results of any sustained dialogue between people in the original position. Rawls claims that he can already show us the results that the participants will necessarily arrive at. His ideal procedure grounds justice not on dialogue but on rational decision.

Moreover, Rawls's veil of ignorance produces agreement by stripping away so much information from the participants that they are for all practical purposes identical. They agree on the principle of maximin—maximizing benefits to the least advantaged—because they have insufficient information about themselves to gain a strategic advantage by refusing to agree. Thus only one rational person is really necessary in the original position, because all rational agents under the veil of ignorance will decide to do the same thing. As a result, not only can we not call the decision a result of dialogue, we cannot even call it an agreement. It is indistinguishable from a single individual's decision of instrumental rationality. This is the ultimate consequence of Rawls's attempt to convert questions of justice into questions of rational decisionmaking. See T. K. Seung, Intuition and Construction: The Foundation of Normative Theory (New Haven: Yale University Press, 1993), 17.

- 14. Thus there is an analogy to the earlier criticism of Rawls. A truly ideal consensus under ideal conditions would require only one ideal participant, because each person in an ideal consensus would know everything (including the perspectives of all of the other parties) and would presumably have the same moral reactions to this knowledge. If the participants did not have the same moral reaction to the same information, it is not clear why they would agree.
- 15. This argument is taken from Balkin, "Transcendental Deconstruction, Transcendent Justice," 1139-40.
 - 16. Ibid.

8. Cultural Heuristics

- 1. I am indebted to Bruce Ackerman for the insight as well as the term.
- 2. Jon Elster, *Making Sense of Marx* (Cambridge: Cambridge University Press, 1985), 466; Leon Festinger, *A Theory of Cognitive Dissonance* (Stanford: Stanford University Press, 1957); Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment Under Uncertainty: Heuristics and Biases* (Cambridge: Cambridge University Press, 1982).
- 3. Elster, Making Sense of Marx, 460-61; Jon Elster, Sour Grapes: Studies in the Subversion of Rationality (Cambridge: Cambridge University Press, 1983), 142.
 - 4. Elster, Sour Grapes, 142.
- 5. Ibid., 141. Elster traces this distinction back to R. P. Abelson, "Computer Simulation of Hot Cognition," in S. Tomkins and S. Messick, eds., *Computer Simulation of Personality* (New York: Wiley, 1963), 277–98.

- 6. Elster, Making Sense of Marx, 466-67; Elster, Sour Grapes, 141.
- 7. See, e.g., Anthony Greenwald and David L. Ronis, "Twenty Years of Cognitive Dissonance: Case Study of the Evolution of a Theory," Psychological Review 85 (1978): 53-57.
- 8. See, e.g., Elliot Aronson, "The Theory of Cognitive Dissonance: A Current Perspective," Advances in Experimental Social Psychology 4 (1969): 1-34, at 16-17; J. Richard Eiser, Social Psychology: Attitudes, Cognition, and Social Behavior (Cambridge: Cambridge University Press, 1986), 93.
- 9. Daniel Kahneman and Amos Tversky, "Judgment Under Uncertainty: Heuristics and Biases," rpt. in Judgment Under Uncertainty, 1-20, at 1; R. Nisbett and L. Ross, Human Inference: Strategies and Shortcomings of Social Judgment (Englewood Cliffs, N.J.: Prentice-Hall, 1980), 6-7.
 - 10. Elster, Sour Grapes, 164-65; Elster, Making Sense of Marx, 505.
- 11. Claude M. Steele and Thomas J. Liu, "Dissonance Processes as Self Affirmation," Journal of Personality and Social Psychology 45 (1983): 5-19; Ruth Thibodeau and Elliot Aronson, "Taking a Closer Look: Reasserting the Role of the Self-Concept in Dissonance Theory," Personality and Social Psychology Bulletin 18 (1992): 591-601; Greenwald and Ronis, "Twenty Years of Cognitive Dissonance," 55; Aronson, "The Theory of Cognitive Dissonance," 27.
 - 12. Elster, Sour Grapes, 148.
 - 13. Ibid., 156.
- 14. Elster, Making Sense of Marx, 510; the quotation is from Karl Marx, "Contribution to the Critique of Hegel's Philosophy of Right: Introduction," in The Marx-Engels Reader, Robert C. Tucker, ed. (New York: Norton, 1972), 11-23, at 12.
 - 15. Elster, Making Sense of Marx, 482.
 - 16. Ibid.
 - 17. Ibid., 466.
- 18. On sample bias and availability heuristics, see Nisbett and Ross, Human Inference, 77-89; Kahneman, Slovic, and Tversky, Judgment Under Uncertainty, 163-208.
- 19. See Amos Tversky and Daniel Kahneman, "The Framing of Decisions and the Rationality of Choice," Science 211 (1981): 543-58; Elster, Making Sense of Marx, 466.
 - 20. Elster, Sour Grapes, 144.
- 21. Karl Mannheim, Ideology and Utopia: An Introduction to the Sociology of Knowledge (New York: Harcourt Brace Jovanovitch, 1936), 118-46.
 - 22. Elster, Making Sense of Marx, 490.
- 23. Ibid., 487. Nevertheless, Elster also notes that "the exploiting classes can be victims of similar illusions. Cognitively based ideologies do not always operate to the benefit of the ruling classes."
- 24. Ibid., 488. Elster draws here on Paul Veyne's work. The basic argument is that "since I would be worse off without a master, it follows on this logic that a society without masters would be intolerable, for who would then provide employment and protection?"
 - 25. See ibid., 322.

- 26. See, e.g., ibid., 464–65, 468–72. Although his discussion focuses almost exclusively on effects of class interests and class positions, it is interesting to note that his actual definition of ideology does not specifically refer to economic class.
- 27. Paul Ricoeur, Lectures on Ideology and Utopia, George H. Taylor, ed. (New York: Columbia University Press, 1986), 8–10, 156–58. Thus, according to Ricoeur, ideology distorts "praxis as something symbolically mediated" (157). This argument provides yet another reason to abandon the familiar base-superstructure model of ideology inherited from Marxism. What that metaphor places in the so-called superstructure (culture) is actually basic to human existence and meaningful human action. Moreover, the superstructure does not exist purely for the purpose of distortion; it is not exhausted by its distorting effects. For example, Ricoeur argues that capitalist understandings of wage labor involve a distortion of praxis because the juridical concept of contract is applied to a situation of domination. But this does not mean that the idea of a contract is merely a fantasy or wholly an element of distortion. Rather, this tool of understanding is more than its distorting effects; it has independent uses, functions, and consequences. Nevertheless, it has been applied to a situation to which it is not fully appropriate; hence it gives a social situation an air of legitimacy that it does not deserve (155–56).
- 28. Clifford Geertz, "Ideology as a Cultural System," in *The Interpretation of Cultures* (New York: Basic, 1973), 209–13. In Chapter 11 we will consider metaphor once again through the work of George Lakoff and Mark Johnson, who argue that these rhetorical devices arise through a process of evolutionary development originating in the movements of the human body.
 - 29. Ibid., 211.
 - 30. Ibid., 212.
- 31. See George Lakoff and Mark Turner, More Than Cool Reason: A Field Guide to Poetic Metaphor (Chicago: University of Chicago Press, 1989).
- 32. Elster, Making Sense of Marx, 492-93. See Karl Marx, The Eighteenth Brumaire of Louis Bonaparte, in The Marx-Engels Reader, 436-525.
 - 33. Marx, The Eighteenth Brumaire, 437.
- 34. Aristotle, On Rhetoric: A Theory of Civic Discourse, George A. Kennedy, trans. (Oxford: Oxford University Press, 1991).
- 35. George Herbert Mead, Mind, Self, and Society, Charles W. Morris, ed. (Chicago: University of Chicago Press, 1962), 154–56, 178–226.
- Jerome A. Bruner, Acts of Meaning (Cambridge: Harvard University Press, 1990),
 138.

9. Narrative Expectations

- Jerome A. Bruner, Acts of Meaning (Cambridge: Harvard University Press, 1990),
 47.
- 2. Erving Goffman, Frame Analysis: An Essay on the Organization of Experience (New York: Harper and Row, 1974).
- 3. Roger C. Schank and Robert Abelson, *Scripts, Plans, Goals, and Understanding* (Hillsdale, N.J.: Erlbaum, 1977).

- 4. Paul Grice, Studies in the Way of Words (Cambridge: Harvard University Press, 1989), 26-27.
 - 5. Bruner, Acts of Meaning, 49-50.
- 6. Roger C. Schank, The Connoisseur's Guide to the Mind: How We Think, How We Learn, and What It Means to Be Intelligent (New York: Summit, 1991), 187-99.
- 7. Roger C. Schank, Tell Me a Story: A New Look at Real and Artificial Memory (New York: Scribner's, 1990), 114-15.
- 8. See Jean Matter Mandler, Stories, Scripts, and Scenes: Aspects of Schema Theory (Hillsdale, N.J.: Erlbaum, 1984).
- 9. F. C. Bartlett, Remembering: A Study in Experimental and Social Psychology (Cambridge: Cambridge University Press, 1932), 63-94.
- 10. Schank, Tell Me a Story, 46; Elizabeth Stone, Black Sheep and Kissing Cousins: How Our Family Stories Shape Us (New York: Times Books, 1988), 96-108, 165-95.
- 11. Donald Polkinghorne, Narrative Knowing and the Human Sciences (Albany: SUNY Press, 1988), 150.
 - 12. Bruner, Acts of Meaning, 121.
 - 13. Ibid., 121-22.
- 14. Donald Spence, Narrative Truth and Historical Truth: Meaning and Interpretation in Psychoanalysis (New York: Norton, 1984); Roy Schafer, "Narration in the Psychoanalytic Dialogue," in On Narrative, W. J. T. Mitchell, ed. (Chicago: University of Chicago Press, 1981).
- 15. Mark J. Osiel, "Ever Again: Legal Remembrance of Administrative Massacre," University of Pennsylvania Law Review 144 (1995): 463-704, at 475-77.
 - 16. Ibid., 476.
- 17. Richard Slotkin, Gunfighter Nation: The Myth of the Frontier in Twentieth-Century America (New York: Atheneum, 1992).
 - 18. Ibid., 10-12.
 - 19. Ibid., 14.
 - 20. Ibid., 16-18.
 - 21. Ibid., 10-21.
 - 22. Ibid., 13.
- 23. Indeed, the story of Exodus contains a cycle within a cycle: The Jews are almost within sight of Canaan and are about to receive the Law when they stray from God's teachings and worship the Golden Calf. As a result, God punishes them by condemning them to wander forty years in the desert before permitting them to enter the Promised Land.
- 24. As Holmes pointed out, even a dog understands the difference between these two states of affairs. Oliver Wendell Holmes, Jr., The Common Law (1888), Mark DeWolfe Howe ed. (Boston: Little, Brown, 1963), 7.
 - 25. See Osiel, "Ever Again," 470 n. 23.
- 26. See Martha Minow, Making All the Difference: Inclusion, Exclusion, and American Law (Ithaca, N.Y.: Cornell University Press, 1990).
 - 27. Ibid.

10. Homologies and Associations

- 1. See, e.g., Kimberlé Williams Crenshaw, "Race, Reform, and Retrenchment: Transformation and Legitimation in Antidiscrimination Law," Harvard Law Review 101 (1988): 1331-87, at 1370-74.
- 2. See Office of Applied Studies, U.S. Department of Health and Human Services, Preliminary Estimates from the 1994 National Household Survey on Druge Abuse (Advance Report Number 10, 1995), 76 (estimating that 76 percent of illicit drug users are white, compared with 13.67 percent who are black); Martina Shea, U.S. Department of Commerce, Dynamics of Economic Well-Being: Program Participation, 1990-1992 (1995), 13 (reporting that 64.5 percent of recipients in major means-tested aid programs are white, while 31.3 percent of such recipients are black); Kathleen Maguire and Ann L. Pastore, eds., Bureau of Justice Statistics, U.S. Department of Justice, Sourcebook of Criminal Justice Statistics (1995), 408 (reporting that 66.6 percent of all criminal arrestees in 1994 were white, while only 31.3 percent of arrestees were black). Although it is true that the percentages of blacks in these groups is higher than their representation in the U.S. population, these statistics nevertheless undercut the all too prevalent assumption that blacks are responsible for the vast majority of these social problems.
- 3. See Patricia A. Williams, The Alchemy of Race and Rights: The Diary of a Law Professor (Cambridge: Harvard University Press, 1991).
 - 4. Roland Barthes, Mythologies (New York: Hill and Wang, 1972).
- 5. See, e.g., Robert Hodge and Gunther Kress, Social Semiotics (Ithaca, N.Y.: Cornell University Press, 1988).
- 6. Gerald Torres and Donald Brewster, "Judges and Juries: Separate Moments in the Same Phenomenon," Law and Inequality 4 (1986): 171-88.
- 7. Thus, although structuralism often claims to look for the underlying "grammar" of culture, the analogy is partly misleading. Structuralist and semiotic analyses do not so much reveal a set of articulable rules consciously used by subjects as the results of practical cognitive work that employs networks of association. Semiotic analysis studies what individuals produce using generative tools of understanding rather than articulable "instructions" of cultural software itself.

This distinction is central to Pierre Bourdieu's critique of structuralism. See Pierre Bourdieu, Outline of a Theory of Practice (Cambridge: Cambridge University Press, 1977); Pierre Bourdieu, The Logic of Practice (Cambridge: Polity, 1990). In Bourdieu's words, we must not confuse the opus operandum with the modus operandi; that is, we must not confuse the structural features of the cultural artifacts produced with the cognitive capacities that produced them. Our conceptual tools create cultural artifacts with a particular conceptual structure. The thing produced is evidence of the generative capacities that produced it; it is not, however, identical with those capacities. Rather than the compiler of a social grammar or a social etiquette book, the semiotician should think of herself as a forensic scientist who, presented with the victim's body, tries to imagine the murder weapon and the means and opportunity for the crime; or as a geologist who, presented with the sedimentary layers of culture, tries to discern the processes that led to that sedimentation.

8. J. M. Balkin, "Nested Oppositions," Yale Law Journal 99 (1990): 1669-1705.

- 9. The standard of appellate review of a jury's verdict also reflects the opposition between law and fact. Questions of law are reviewed de novo; that is, the higher court judge may substitute her own view of the law for that of the lower court judge. On the other hand, an appellate court is loath to reverse factual findings by a jury unless they are completely unreasonable, or, in the case of a trial by a lower court judge, clearly erroneous. This suggests a further homology between appellate court and trial court judges, on the one hand, and judges and juries on the other. Appellate courts are usually called "higher" courts, trial courts are usually called "lower" courts. One appeals to a higher court from a lower court. The work of the higher courts is concerned almost exclusively with matters of law.
 - 10. Torres and Brewster, "Judges and Juries," 181.
- 11. The study of conceptual oppositions does not exhaust the nature of patriarchal thought. As noted in Chapter 5, patriarchy is the result of a jumble of variegated and heterogenous mechanisms. Here I focus on only one of the many different ideological mechanisms of patriarchy.
- 12. See, e.g., Catharine MacKinnon, "Difference and Dominance," in Feminism Unmodified: Discourses on Life and Law (Cambridge: Harvard University Press, 1987), 32-
- 13. Jeanne L. Schroeder, "Abduction from the Seraglio: Feminist Methodologies and the Logic of Imagination," Texas Law Review 70 (1991): 109-210.
- 14. Stephen L. Carter, Reflections of an Affirmative Action Baby (New York: Basic, 1991), 47-69.
 - 15. Ibid., 50-52.
- 16. Andrew Koppelman, Antidiscrimination Law and Social Equality (New Haven: Yale University Press, 1996), 146-76; Silvia A. Law, "Homosexuality and the Social Meaning of Gender," Wisconsin Law Review 1988: 187-235; Marc A. Fajer, "Can Two Real Men Eat Quiche Together? Storytelling, Gender-Role Stereotypes, and Legal Protection for Lesbians and Gay Men," University of Miami Law Review 46 (1992): 511-651.
- 17. J. M. Balkin, "The Constitution of Status," Yale Law Journal 106 (1997): 2313-74, at 2361-62.
 - 18. Fajer, "Can Two Real Men Eat Quiche Together?" 607-29.
 - 19. Balkin, "Nested Oppositions."
 - 20. Ibid., 1676.
- 21. On the Western tradition, see G. E. R. Lloyd, Polarity and Analogy: Two Types of Argumentation in Early Greek Thought (Cambridge: Cambridge University Press, 1966); on the Eastern tradition, see Lao-Tzu, Tao Te Ching, Thomas Cleary, trans. (London: Shambala, 1983).
- 22. See Jacques Derrida, "Différance," in Margins of Philosophy (Chicago: University of Chicago Press, 1987), 1-27. Derrida describes his argument in terms of presence and absence of concepts. He argues that "the signified concept is never present in and of itself, in a sufficient presence that would refer only to itself" (11). A concept's lack of full self-presence is its conceptual dependence upon other concepts. Derrida's "différance" also adds a temporal dimension to the notion of conceptual dependence, for he argues that concepts depend on their predecessors and their successors.

- 23. These claims are defended at greater length in J. M. Balkin, "Transcendental Deconstruction, Transcendent Justice," Michigan Law Review 92 (1994): 1131-86.
- 24. Akhil Amar, "The Bill of Rights as a Constitution," Yale Law Journal 100 (1991): 1131-1210, at 1195.
- 25. Amar points out that the constitutional role of juries as populist protectors of human rights was not simply to nullify in sympathetic factual cases, but to apply the higher law of the constitution to control the ordinary positive law. Ibid., 1191. This suggests the homology reason: passion:: jury: judge:: higher law: positive law:: controlling: needing to be controlled. Indeed, Amar notes, "If we seek a paradigmatic image underlying the Bill of Rights, we cannot go far wrong in picking the jury" (1190). This role, he argues, has been suppressed over time, as judges have assumed for themselves the mantle of protectors of constitutional rights, regarding the jury as unlearned and thus as a threat to constitutional values. See 1192-95.
- 26. As argued by former Supreme Court Justice William Brennan in William Brennan, "Reason, Passion, and 'The Progress of the Law,' " Cardozo Law Review 10 (1988): 3-24.
- 27. Oliver Wendell Holmes, Jr., "The Path of the Law," Harvard Law Review 10 (1897): 457-78.
- 28. For two quite different versions of this point see Andrew Ortony, Gerald L. Clare, and Allan Collins, The Cognitive Structure of Emotions (Cambridge: Cambridge University Press, 1988); and Robert H. Frank, Passions Within Reason: The Strategic Role of the Emotions (New York: Norton, 1988).
- 29. An example would be the medical profession's opinions on proper medical procedures for women.

11. Metaphor, Metonymy, and Cognitive Models

- 1. See Raymond W. Gibbs, Jr., "Process and Products in Making Sense of Tropes," in Metaphor and Thought, Andrew Ortony, ed. (Cambridge: Cambridge University Press, 2d ed., 1994), 252-76.
- 2. George Lakoff, Women, Fire, and Dangerous Things: What Categories Reveal About the Mind (Chicago: Chicago University Press, 1987); Mark Johnson, The Body in the Mind: The Bodily Basis of Reason and Imagination (Chicago: University of Chicago Press, 1987).
 - 3. Lakoff, Women, Fire, and Dangerous Things, 271-78.
- 4. For an attempt to trace these ideas through many different fields and traditions, including Eastern religious thought, see Francisco J. Varela, Evan Thompson, and Eleanor Rosch, The Embodied Mind: Cognitive Science and Human Experience (Cambridge: MIT Press, 1991). For a discussion in the European phenomenological tradition see Hubert L. Dreyfus, What Computers Can't Do: The Limits of Artificial Intelligence (New York: Harper and Row, rev. ed., 1979), 235-55.
 - 5. Pierre Bourdieu, The Logic of Practice (Cambridge: Polity, 1990), 68-79.
- 6. Giambattista Vico, The New Science, Book I [237], Thomas Goddard Bergin and Max Harold Fisch, trans. (Ithaca, N.Y.: Cornell University Press, 1968), 78.

- 7. George Lakoff, "The Contemporary Theory of Metaphor," in Metaphor and Thought, 202-51.
- 8. George Lakoff and Mark Johnson, Metaphors We Live By (Chicago: University of Chicago Press, 1980), 90, 98.
 - 9. Ibid., 98-99.
 - 10. Ibid., 87-96.
- 11. Michael J. Reddy, "The Conduit Metaphor: A Case of Frame Conflict in Our Language About Language," in Metaphor and Thought, 164-201.
- 12. See Eve Sweetser, From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure (Cambridge: Cambridge University Press, 1990).
- 13. Steven Winter, for example, has shown how metaphorical models of understanding have shaped the growth and development of legal doctrine. Steven L. Winter, "The Meaning of 'Under Color of Law," Michigan Law Review 91 (1992): 323-418; Steven L. Winter, "Transcendental Nonsense, Metaphoric Reasoning, and the Cognitive Stakes for Law," University of Pennsylvania Law Review 137 (1989): 1105-1237; Steven L. Winter, "The Metaphor of Standing and the Problem of Self-Governance," Stanford Law Review 40 (1988): 1371-1515.
 - 14. See Lakoff and Johnson, Metaphors We Live By, 4-5, 77-87.
- 15. The locus classicus of this metaphor in American free speech law is Justice Holmes's dissent in Abrams v. United States, 250 U.S. 616 (1919).
 - 16. Lakoff and Johnson, Metaphors We Live By, 157.
 - 17. Ibid., 5.
 - 18. Ibid., 4.
- 19. Donald A. Schon, "Generative Metaphor: A Perspective on Problem-Setting in Social Policy," in Metaphor and Thought, 137-63.
- 20. Lakoff and Johnson, Metaphors We Live By, 35-40; Gibbs, "Process and Products in Making Sense of Tropes," in Metaphor and Thought, 252-76, at 258-62.
 - 21. Lakoff, Women, Fire, and Dangerous Things, 78.
- 22. Ibid., 16, 161. Douglas Medin, "Concepts and Conceptual Structure," American Psychologist 44 (1989): 1469-81, at 1470.
- 23. Ludwig Wittgenstein, Philosophical Investigations (Oxford: Blackwell, 1953), 66-71.
- 24. Eleanor Rosch, "Cognitive Reference Points," Cognitive Psychology 7 (1975): 532-47; Eleanor Rosch, "Cognitive Representations of Semantic Categories, Journal of Experimental Psychology: General 104 (1975): 192-233; Eleanor Rosch and C. B. Mervis, "Family Resemblances: Studies in the Internal Structure of Categories," Cognitive Psychology 16 (1975): 371-416; Eleanor Rosch and B. B. Lloyd, eds., Cognition and Categorization (Hillsdale, N.J.: Erlbaum, 1978).
- 25. Lance J. Ripps, "Inductive Judgments About Natural Categories," Journal of Verbal Learning and Verbal Behavior 14 (1975): 665-81; Rosch, "Cognitive Reference Points"; Amos Tversky and I. Gati, "Studies of Similarity," in Rosch and Lloyd, Cognition and Categorization, 79-88; Lakoff, Women, Fire, and Dangerous Things, 41.
 - 26. Lakoff, Women, Fire, and Dangerous Things, 91.
 - 27. Ibid., 43–44, 68.

- 28. For a survey of the literature and the relevant debates, see Medin, "Concepts and Category Formation." A different account is offered in Lakoff, Women, Fire, and Dangerous Things.
 - 29. Lakoff, Women, Fire, and Dangerous Things, 85.
- 30. Amos Tversky and Daniel Kahneman, "Probability, Representativeness, and the Conjunction Fallacy," *Psychological Review* 90 (1983): 293–315.
 - 31. Lakoff, Women, Fire, and Dangerous Things, 85-86.
 - 32. Ibid., 80-84.
 - 33. Ibid.
 - 34. Cf. ibid., 81.
 - 35. Ibid.

12. The Power of Understanding

- 1. Michel Foucault, "Truth and Power," in *Power/Knowledge: Selected Interviews and Writings*, 1972–1977, Colin Gordon, ed. (New York: Pantheon, 1980), 118.
 - 2. Ibid.
- 3. See, e.g., Michel Foucault, Discipline and Punish: The Birth of the Prison (New York: Vintage, 1979), 170.
 - 4. Foucault, "Truth and Power," 118.
 - 5. See Foucault, "Truth and Power," 131.
- 6. Charles Taylor, "Foucault on Freedom and Truth," in Foucault: A Critical Reader, David Couzens Hoy, ed. (Oxford: Blackwell, 1986), 69–102. Ironically, although the goal of genealogy is to deny that a deeper hidden meaning lurks beneath the surface of social events, this truth itself must be revealed through a process of unmasking the fraud of deep meaning. The deeper meaning of social life is that there is no deeper meaning. Dreyfus and Rabinow demonstrate this paradox in their very formulation of Foucault's project: "The genealogist recognizes that the deep hidden meanings, the unreachable heights of truth, the murky interiors of consciousness are all shams... Genealogy's... banner [is]: Mistrust identities in history: they are only masks." Hubert L. Dreyfus and Paul Rabinow, Michel Foucault: Beyond Structuralism and Hermeneutics (Chicago: University of Chicago Press, 2d ed., 1983), 107. Yet mistrust implies a truer state of affairs that lies behind what is mistrusted.
- 7. See Nicos Poulantzas, State, Power, and Socialism, trans. Patrick Camiller (London: New Left, 1978), 149.
- 8. Michel Foucault, *The History of Sexuality*, vol. 1, *An Introduction* (New York: Vintage, 1980), 95–96.
 - 9. Ibid., 94.
- 10. Thomas Seung has suggested to me that Foucault's theory of resistance is Hegelian, because the antithesis (resistance) grows magically out of the thesis (the system of power).
 - 11. See Foucault, Power/Knowledge, 82-83.
 - 12. Ibid., 81.
 - 13. See, e.g., Foucault, The History of Sexuality, 1: 95.

- 14. See Michel Foucault, "Nietzsche, Genealogy, History," in Language, Counter-Memory, Practice: Selected Essays and Interviews by Michel Foucault (Ithaca, N.Y.: Cornell University Press, 1977), 148–51.
- 15. See, e.g., Foucault, "Nietzsche, Genealogy, History," 147-48; Discipline and Punish, 25-30.
 - 16. Foucault, The History of Sexuality, 1: 94.
- 17. Michel Foucault, *The History of Sexuality*, vol. 2, *The Use of Pleasure* (New York: Vintage, 1985), 6–7.
 - 18. Ibid., 7.
 - 19. Ibid., 6.
 - 20. Foucault, "Nietzsche, Genealogy, History," 142.
- 21. Stanley Fish, *Doing What Comes Naturally: Rhetoric and Change in Law and Literary Studies* (Durham, N.C.: Duke University Press, 1990), 520 (italics omitted).
 - 22. See Charles Taylor, "Foucault on Freedom and Truth."
- 23. See Stephen Jay Gould, *Hen's Teeth and Horse's Toes* (New York: Norton, 1983), 158-65.
- 24. Hans-Georg Gadamer, Truth and Method (New York: Crossroad, 1975), 238-40.
- 25. Stanley Fish, There's No Such Thing as Free Speech (and It's a Good Thing Too) (Oxford: Oxford University Press, 1994), 117.
 - 26. Foucault, The History of Sexuality, 1: 94-95.
 - 27. Gadamer, Truth and Method, 266-67, 324-25.
- 28. See Michel Foucault, "An Aesthetics of Existence," in *Foucault Live: Interviews*, 1966–84 (New York: Semiotext(e), 1989), 313.
- 29. Like all articulation, this process involves construction as well as refinement; thus Foucault is partly correct that articulation does not involve the rediscovery of a deeper sexual nature that was always present. That is because sexual desire, like all human desires and values, is inchoate and indeterminate. It must be articulated through the development of culture. Although sexual desire is articulated through culture, sexual desire is not wholly a creation of culture; even before culture existed, human beings had sexual desires. Foucault is ambiguous on this point. He doubts that "sex is an anchorage point that supports the manifestations of sexuality"; rather, he thinks it is "a complex idea that was formed inside the deployment of sexuality." Foucault, The History of Sexuality, 1: 152. Sex is a concept that we use to describe the ways in which we have understood our bodies through culture. Thus, Foucault insists, "sex is not an autonomous agency which secondarily produces manifold effects of sexuality over the entire length of its surface of contact with power." Instead, "sex is the most speculative, most ideal, and most internal element in a deployment of sexuality organized by power in its grip on bodies and their materiality, their forces, energies, sensations, and pleasures" (155).

The difficulty with this formulation lies in the last phrase. How can power have a grip on the "energies, sensations, and pleasures" of bodies if bodies have energies, sensations, and pleasures only as a result of culture? Here Foucault's Parmenideanism reasserts itself: sex must always have been internal to sexuality; everything must already be

fully contained within the system of cultural power. Yet without human values to be shaped through culture, cultural articulation cannot even get off the ground.

30. Immanuel Kant, "Conjectural Beginning of Human History," in Kant on History, Lewis White Beck, ed. (New York: Macmillan, 1963), 57.

13. Knowledge Made Flesh

1. Jerome H. Barkow, Leda Cosmides, and John Tooby, The Adapted Mind: Evolutionary Psychology and the Generation of Culture (Oxford: Oxford University Press, 1992).



Adorno, Theodore, 26 African Americans, 112, 115, 117-18; associations with, 254; conceptual oppositions relating to, 232-33. See also racism Amar, Akhil, 237, 322n.25 American Revolution: meaning of, 212 animal rights, 146-47 antihumanism, 276-77 anti-Semitism, 106, 108, 112; Japanese, appellate courts, 321n.9 argument: metaphoric models for, 244, 245-48 Aristotle, 183 art: as method of memetic transfer, 194 articulation of human values, 27-31; as source of human freedom, 284-85; theories of justice as, 120, 314n.4 artificial intelligence, 4 artificiality, 29 autobiographical narratives, 200 autonomy: and hermeneutic power, 280 availability heuristic, 178-79, 182, 253

bandwagon effect, 84
Barrett, Michéle, 103
Barry, Marion, 254
Barthes, Roland, 217, 263
behavior: effect of memes on, 66–68, 305n.63; as explanation of shared conventions, 10–11; and narrative expectations, 213
behaviorism, 10, 186

beliefs: and cultural software, 44; as filters, 58; memes as, 44, 76-77, 308n.31; propositional, 102-3; reflexive, 82-83; survival of, 64 "Best Black" syndrome, 232 Bible: story of Exodus, 206-7, 319n.23; story of Joseph, 81, 307n.24 biological bricolage, 88-90 biological evolution, 48-49, 64-65, 290; analogy with cultural bricolage, 37-41; as distinguished from cultural evolution, 35-37; and memetic evolution, 48 black Americans. See African Americans bodily movement, 43; metaphors based on, 242-43 body: in Foucault, 267-68 Bourdieu, Pierre, 33-34, 186, 243, 263, 320n.7 brain structure: and memes, 45-46; uniqueness of, 5-6 Brewster, Donald, 219, 224, 225 bricolage, xi. See also biological bricolage; conceptual bricolage; cultural bricolage; metabricolage Bruner, Jerome, 4, 185

The Care of the Self (Foucault), 283
Carter, Stephen, 232
Cassirer, Ernst, 310nn.2, 3
categories: as necessary for understanding information, 18–19, 77; as nested oppositions, 234–35; and prototypes, 250–58
categorization: cognitive models of, 250–58

causation, problem of, 9 Chomsky, Noam, 12, 13; and innate ideas, 217, 218 Civil War, 204 class, 123, 178-80; Marxist analyses of, 133. See also Marxist theories of ideology cognitive biases, 176, 178-79, 183 cognitive dissonance, 52, 58, 102, 175, 176-77, 184, 258, 278 cognitive mechanisms, 58, 102 cognitive models, 242-43; and categorization, 250-58; metaphoric, 243-45; metonymic, 248-49 cognitive science, 4 collective consciousness: Durkheim's concept of, 10, 95-96, 268 collective memory, 203-4 commensalism, 62, 63, 70 communication: as factor in spread of memes, 81-82, 90-92; metaphors for, 244; and possibility of manipulation, 275; process of, 300-301n.13 comprehension: as factor affecting survival of memes, 79–81 computer: and memetic growth, 55-57; as metaphor, 4-6, 14-17, 92-94, 134 conceptual bricolage, 92, 104 conceptual homologies, 219-21; hermeneutic function of, 221; and hierarchies, 223-26 conceptual oppositions, 216; clusters of, 223-26; compared with cognitive models, 245, 257-58; hierarchies of, 223-30; and human freedom, 52, 279-80, 284-85, 293-94; judge and jury as, 219-23, 224, 322n.25; mediations between, 221-23; memetic approach to, 218-19; nested, 222-23, 234-35; privileging of elements of, 230-34; and structuralism, 217-19; subcategorization within, 222; suppression and projection in, 235 - 41Conjectural Beginning of Human History (Kant), 284 conventions, shared: criticism of social coordination model of, 11-12, 296-97n.21; explanations for, 10-14; memes as building blocks of, 68-72 copyright law, 47-48, 301-2n.18 creativity, 52, 163; and narrative structures, 194-96

cultural associations: ideological effect of, 216-17, 218. See also conceptual oppositions cultural bricolage: analogy with biological evolution, 37-41; cultural software as result of, xi, 31-35; economy of, 33-34; examples of, 33; as a form of cultural evolution, 35-37, 160; and historical development, 39-40; imperfections of, 40; and memetic evolution, 72-73; and theories of evolution, 35-37; unintended effects of, 34-35 cultural constructions, 28-29 cultural evolution: attraction model of, 303n.32; and biological evolution, 35-37, 48; combining of lineages in, 52-53; Darwinian mechanism at work in, 42, 303n.32; "demand" side versus "supply" side effects in, 54; Lamarckian evolution, 35-37; memes as unit of inheritance in, 43-48, 55; and symbolic forms, 102; systematic bias in transformations required for, 53-54. See also memes cultural hegemony, 68 cultural heuristics, 180-87; narratives as, 208-9 cultural know-how, 6-7, 287; and narrative structures, 193. See also cultural software cultural proliferation, 281-85 cultural relativism, 150-55 cultural separation, 91 cultural software: and agency, 52, 279-80, 284-85, 293-94; as alternative to power/ knowledge, 272-73; ambivalent nature of, 126-29, 141, 152-53, 248; and conceptual oppositions, 218-19; as constituting human existence in history, 5-6, 17-18, 287-88; critical approach to, 130-38; cumulative development of, 32; defined, ix, 3, 6; distribution of, 96–97; economy of, 92-96, 136, 277-78; effects on behavior of, 67-68; as empowering force, 16-17, 18; endemic forms of, 61; epidemic forms of, 61; as essential to reason, 138-41, 288-94; evolution of, 5, 95-96, 286-87; as explanation for shared understandings, 13-14; of a group, 94; hermeneutic power of, 270–71, 273–76; ideological effects of, 104-5; of individuals, 93-94; as informational

virus, 110; and Mannheim's paradox, 128;

as means of cultural transmission, x-xi, 5-6; memes as building blocks of, 43-48; as metaphor, 4-6, 14-17, 92-94, 287; and power/knowledge, 269-72; power of, 274; reproduction of, 90-92; as result of bricolage, 31-35; self-referential nature of, 128-29; semiotic systems as, 217-19; survival of, 53-54; as tool, 3-4, 6-7, 14; transformations of, 53-54; values transmitted by, 27-31. See also conceptual oppositions; ideology; memes; narrative structures

cultural software, theory of: as distinguished from Foucault's theory of power/knowledge, 272-73; and theories of cultural understanding, 7-14; as theory of historical existence, 17-18; as theory of informational existence, ix, 14-15, 287-88

cultural tools, 24-27; as expression of human values, 27-31; heuristic as, 139, 140-41; language as, 24; metaphor as, 139, 140-41; unintended uses of, 32-33; uses of, 24-26

cultural understanding: change as aspect of, 278-79; and dynamic nature of traditions, 95-96; Gadamer's theory of, 7-10, 13, 14, 95, 313n.18; and lines of memetic descent, 88-90; tools of, 1-2, 3, 24-26, 38-39, 45, 96-97, 104, 108, 126-27, 128-29, 139-41, 286. See also cultural software; narrative structures

culture: as bricolage, xi, 31-35; cumulative power of, 5; Foucault's view of, 267; intersubjective aspects of, 26-27; memetic development of, 142-43; philosophy of, 2-3; as system of inheritance, 43; theory of, x; transmission of, x-xi, 5-6, 43-45, 46, 47, 51, 53, 92–94, 102

Darwin, Charles, 35-37, 52 Darwinian algorithms, 290 Darwinian evolution, 35-37, 41, 42, 298n.16 Davidson, Donald, 312-13n.14 Dawkins, Richard, 43, 72 deconstruction: as theory of nested oppositions, 234–35 default characterizations, 253-54

defense mechanisms, 66, 75 Delius, Juan, 302n.22 Dennett, Daniel, 46, 47, 55, 57, 60, 90, 301n.16 Derrida, Jacques, 234, 235, 321n.22 designed objects, 37 designoid objects, 37 dharmatzedek, 163-67 dialogic encounters, 148-49 différance, 235, 321n.22 differentiation, problem of, 9-10 dikaiosunē, 165-67 discourse, 2; Foucault's analysis of, 26-27, 267; theory of, 186 disease: as metaphor for ideology, 108-10; as metaphor for memes, 60-61 dissonance reduction, 176-77, 180, 184, 258, 278 domination: justice as consideration in, 118-20; Thompson's definition of, 118 "double whammy," or double penalty, 231-34, 257 Dreyfus, Hubert L., 324n.6 Durkheim, Emile, 10, 96

Eagleton, Terry, 126 ecologies: cultures as, 83-90; and memetic growth, 78-79 education, 66 ego defense, 3. See also cognitive dissonance The Eighteenth Brumaire of Louis Bonaparte (Marx), 182, 183 Ellickson, Robert, 89 Elster, Jon, 101, 104, 107; and microfoundations for Marxist theory, 175-80 equality, 146, 155-56 ethical values, articulation of, 29 etymology, 298n.11; cognitive modeling in, 245; as cultural software, 89; Vico's theory of, 243. See also language evidence, rules of, 220 evolution: by natural selection, 42. See also biological evolution; cultural evolution expectations, 70. See also narrative structures

faith: memes for, 75 fallacy of composition, 179 false consciousness, 75, 116-18, 123 felons, 118-19

females. See women
Festinger, Leon, 175, 176, 258
Fish, Stanley, 135, 271, 273
Fiss, Owen, 313n.1
Forms, Plato's theory of, 144
Foucault, Michel, 26–27, 280; on
discourses of sexuality, 281; theory of
power/knowledge, 261–71, 272–73, 280–
81, 281–82
freedom of speech, 87; as issue in Zenger
trial, 237; and metaphor of "marketplace
of ideas," 246
frontier: as myth in American culture, 205–6

Gadamer, Hans-Georg, 186, 273, 277, 311-12n.14; theory of cultural understanding, 7-10, 13, 14, 95, 313n.18; theory of tradition, 7-10, 13, 14, 263, 297 Gardner, Howard, 4 Geertz, Clifford, 102, 123, 242; conception of ideology, 180 Gellner, Ernest, 140 gender: associations, 34; conceptual oppositions applied to, 228-30; and false consciousness, 117; stereotypes of, 224-26, 233-34 genealogy: Foucault's notion of, 266, 268, 270, 324n.6 genes, 45, 47 genetic transmission, 47 genomes, 46 Goffman, Erving, 191, 192 Gould, Stephen Jay, 36-37, 38, 39-40, 71-Grice, Paul, 192 Guess, Raymond, 123

Habermas, Jürgen, 8, 160, 314n.4
habitus, theory of, 243
Hamlet (Shakespeare): pigeon immune to memes contained in, 55, 196
hardware: rejection of concept of biological, 4–5
Harlan, John, 87
Harman, Gilbert, 315n.6
Hegel, G. W. F., 2, 10, 11, 24, 35
hegemony: and false consciousness, 116–18
Heidegger, Martin, 281
Heraclitus, 18
hermeneutic conformation, 132
hermeneutic co-optation, 132, 313n.18

hermeneutic power, 261; and autonomy, 280; of cultural software, 270-71; economy of, 276-79; individuals as agents of, 279-80; mechanisms of, 273-75; susceptibility to, 277-78; subjection to, 279-85 hero: in American mythology of the Frontier, 206 heuristics, 3, 139, 286; cultural, 180-87; as tools of understanding, 140-41 hierarchies: of conceptual oppositions, 223-26 Hill, Anita, 197 Hispanic communities, 115 historical existence: cultural software as theory of, 17-19, 182-83, 287-88 historicism, 167, 169-70 historiography, 82 history: human versus natural, 6; narrative construction of, 214-15 The History of Sexuality (Foucault), 268-69, 276, 280-81 Holmes, Oliver Wendell, 240 Holocaust, 207 homo faber, 23 homologies, 3. See also conceptual homologies homosexuality: conceptual homologies applied to, 233-34 Horkheimer, Max, 26 human values. See values Husserl, Edmund, 12 ideal dialogic theories, 159-60 ideal observer theories, 157-59

ideal process theories, 157-59 ideas: impact of context on, 87-88 ideological drift, 87-88 ideological effects, 3-4, 104-5, 110, 142; causal explanations of, 175-80, 181-82, 185; of cognitive models, 255–58; of cultural associations, 216-17; defined, 104-5; distributed through memetic transmission, 173-74; of metaphors, 245-48; of narratives, 210-11, 215; produced by language, 185 ideological mechanisms, 104-5, 173, 176; defined, 104-5; suppression as, 239 ideological power, 270-71 ideological thinking, 105, 107-8 ideology: ambivalent conception of, 3, 126-29, 280, 282; analysis of, 129-38, 145,

149-50, 312n.14; and cognitive filters, 76; conceptions of, 3-4; contrasting definitions of, 2, 101–3; critical approach to, 130-38; cultural software and, ix, 3-7, 17-19, 61; defined, 2, 101-3; and epidemic and endemic cultural software, 61; Foucault's view of, 262-69; functional approaches to, 103-4; and Gadamer's theory of tradition, 7-8, 263; hegemonic conceptions of, 116; justice and, 105-6, 116-21; malfunctions in, 108-9; and Mannheim's paradox, 125-28; Marxist theories of, 103, 107, 110, 122-23, 132; neutral conceptions of, 122-26, 180; pejorative conceptions of, 122–26, 180, 283, 295n.3; and philosophy of culture, 2-3; positive conception of, 123; power of, 110-11; role of metaphor in, 181; study of, 1-2, 285; as system of symbols and associations, 180-81; total conception of in Mannheim, 124-25; as worldview, 3. See also justice imperialist universalism, 150-55 individualism: methodological, x, 176, individuals: shared consciousness of, 10-11; uniqueness of, 4-5, 49 information: as embodied by human beings, 14-15, 287-88; filtering of, 59; processing of, 18-19, 79; storage of, 55-56, 79; as unifying concept, ix. See also cultural software institutional authority, 85-86 institutions, 6; memes as building blocks of, 68-72; and spread of memes, 86-88 instrumental rationality, 25-26 intellectual property law, 301-2n.18 isolation: impact of on cultural transmission, 90-92

jazz, 28; pianist as human embodiment of information, 287

Jews, 50; Japanese view of, 106; prejudice against, 112; stock story of, 206–7

Jim Crow laws, 155–56

Johnson, Mark, 242, 247

Judaism: stock story of, 206–7

judges: conceptual oppositions associated with, 219–23, 224, 225, 226–27; nested oppositions applied to, 235–41; privileging of over juries, 230–31

juries: conceptual oppositions associated with, 219-23, 224, 225, 226-27, 322n.25; and narrative structures, 197-98; nested oppositions applied to, 235-41; privileging of judges over, 230-31 justice: as basis for theory of ideology, 105-6, 111-21; conceptions of, 142-43, 163-67; and cultural relativism, 153-55; as human value, 30-31; ideal observer theories of, 158-59; ideal of, 147-48; indeterminacy of boundaries of, 146-47; metaphors of, 161-63; Muse of, 161-63; narrative structures and, 209-10; and neutral conception of ideology, 124; positive norms of, 155-57; subjects and agents of, 146-48, 166, 314n.2; transcendent ideal of, 158, 159-60, 163-67; as transcendent value, 143, 144-49; universal standards of, 150-51

Kahneman, Daniel, 175, 176, 183, 253, 258

Kant, Immanuel, 2, 35, 160; conception of reason as viewed by, 138, 284–85; concept of regulative ideals, 145; concept of transcendental subject, 12, 13; and problem of shared understandings, 12–13

Keynes, John Maynard, 84 King, Martin Luther, 87 Korean Americans, 113 Kulturkampfs, 68

Lakoff, George, 242, 247, 255 Lamarck, Jean-Baptiste de Monet de, 35–37

Lamarckian evolution, 35–37 language: abilities as examples of complexes of meme, 43; as cultural tool, 24, 26; evolution of, 33, 88, 298n.11; and ideological effects, 185; impact of isolation on, 91; and memetic evolution, 55, 57. See also metaphors; metonymy

language games, 154, 186, 263 law. See judges; juries Laws (Plato), 30 Lear, King, 24 Lee, Robert E., 204 Lenin, 123 lesbians: stereotypes of, 233–34 Lévi-Strauss, Claude, 12, 13, 24, 82, 258; concept of bricolage, 31, 270; structuralism of, 217–18, 219
Lewis, David, 11
Lewontin, Richard C., 36–37, 71–72
location, problem of, 8
Lukacs, Georg, 123
Lumsden, Charles J., 65, 67
Lyotard, Jean-François, 154–55

MacKinnon, Catharine, 114

McLuhan, Marshall, 79 maladaption, 108-9 males. See men malfunction, 108-9 manipulation, 275 Mannheim, Karl: conception of ideology, 124-28, 179 Mannheim's paradox, 125-28 markets, 37 Marx, Karl, 2, 35, 40; views of tradition of, 183 Marxism, 133 Marxist theories of ideology, 103, 107, 110, 122-23, 132; Elster's microfoundations for, 175-80; Foucault's objections to, 262-63 Mayr, Ernst, 49, 90 media analysis, 79 mediation: between conceptual oppositions, 221-23, 233 meme pool, 94, 95 memes: as building blocks of cultural software, 43-48, 174, 286-87, 300n.8; communication as factor in spread of, 90-92; competition among, 50-51, 61-64, 66-68, 75-97; cultural bricolage and architectural constraints on, 71-73; in cultural conventions, 68-72; and cultural evolution, 43-48, 303n.32; ecological factors affecting, 74, 83-90, 93; examples of, 43-44; as filters, 57-60; in institutions, 68-72, 86-88; mutation in, 51-52, 66, 90; parasitic, 65-66; as populations, 48-51; problems in transmission of, 54-57; psychological factors affecting, 74, 77-83; replication of, 54-57, 74-90, 95; substantive factors affecting, 74-77; survival of, 74-90; susceptibility to, 70-71; as symbionts, 61-68; symbols as, 47; types of, 43;

variation in, 50, 51-54; as viruses, 60memetic evolution, 48, 61-68, 92; and concept of justice, 142; and cultural bricolage, 71-73; through cultural communication, 109-10. See also memes memetic kinship, 68, 305n.74 memetic phenotype, 45 memomes, 46 memory: collective, 203-4; and cultural transmission, 77-78; and narrative structure, 189, 196-97; social, 203-4; as storage for memes, 50, 77-81; of stories, 54, 78 men: associations, 34; stereotypes of, 228-30, 231-32 metabricolage, xi, 175, 186-87 metaphoric models, 243-45 metaphors, 3, 181, 243; alteration of, 89; based on bodily movements, 242-43; extensions of, 33, 298n.11; ideological effects of, 245-48; of justice, 161-63; Marx's use of, 183; of self-reflection, 134; as tools of understanding, 139, 140–41, metonymic models, 248-59; body as metonym for human being, 267-68; categories as, 252; examples of, 248-49; ideological effects of, 249. See also prototypes metonymy, 3, 181, 248 mind, philosophy of, 4 modernity, 96-97 moral relativism, 315n.6 motherhood, models of, 255-57 Muse, myth of, 163 music: and articulation of aesthetic values, 28-29 mutualism, 62, 63, 70 myths: ideological effects of, 207-8; as vehicles of memory storage, 78; narrative structures as, 204-5; and spread of memes, 82

narrative structures, 3, 188–90; and autobiography, 200; and the canonical, 210–11; as cultural heuristics, 208–9; as form of memory storage, 188–90; of groups, 203–9; as heuristics, 208; and justice, 209–10; in legal contexts, 197–98; myths as transformation of, 82; as

networks of expectations, 190-91; as norms, 191-94; and personal relationships, 198-203; shared, 211-15; and the social world, 211-15; as stock stories, 194-98; transmission of, 194-95; uses of, 189 naturalness, 29 natural selection, 64; evolution by, 42; and units of cultural transmission, 61 nested oppositions, 234-35 nesting, 222-23 Nietzsche, Friedrich Wilhelm, 2, 270 nomos, 2 normative order, 166-67 norms: cultural expectations as, 191; positive, 148; narratives as, 191-94; prototypes as, 254, 255-58; transcendent, 148, 150

Objective Spirit: Hegel's concept of, 10, 11, 268 ontology: problem of, 8, 268 oppositions. See conceptual oppositions; nested oppositions oral cultures, 79 The Origin of Species (Darwin), 52 Other: as applied to ideological analysis, 131; as manifested in myth, 206; projection onto an, 237-38 overgeneralization, 252

panda's thumb, 38, 72, 291 parasitism, 62, 63, 65-66, 70 Parmenides, 265 passion: and reason, 220-21, 222, 223-24, 226-27, 236-41 patriarchy, 114, 224; reinforced by conceptual oppositions, 228-30 Peirce, Charles Sanders, 315-16n.12 physis, 2 Pirke Avot, 1, 128 Plato: and justice, 30, 144, 165; theory of Forms, 144 populations: species as, 49 Poulantzas, Nicos, 264 power: Foucault's theory of, 261-69; resistances to, 264-66. See also hermeneutic power pragmatism, 167-69 prejudice, 114. See also ideological effects print media, 80

problematization: in Foucault, 283-84 projection, 237 prototypes, 250-52; asymmetrical inferences from, 254; as default characterizations, 253-54; as evaluative norms, 255-58; as indicators of relative prevalence, 253; as reference points, 254 psychotherapy: and personal scripts, 201-3

Rabinow, Paul, 324n.6

race: viewed through conceptual oppositions, 216, 218 racism, 108, 113; as epidemic and endemic, 61; factors contributing to, 258; and hegemony model of ideology, 116; as manifestation of cultural knowledge, 211; as viewed by competing ideologies, 124. See also Jim Crow laws; ideological effects rational-choice theorists, 70 Rawls, John, 316n13 Reagan, Ronald, 253 reason: as applied to cultural software, 288-94; as cultural heuristic, 138-41; evolution of, 138; idea of, 240; Kant's view of, 138, 284-85; and memetic evolution, 288-94; reflexive nature of, 139; and passion, 220-21, 222, 223-24, 226-27, 236-41 regulative ideals, 145 religion, 75; Marx's view of, 177-78 representations: as units of cultural transmission, 44 Republic (Plato), 30, 165 resistance, Foucault's theory of, 264-66 rhetoric, 19; as ideological mechanism, 183; "just," 140-41 Ricoeur, Paul: on ideology, 180-81, 318n.27 Rosch, Eleanor, 250

salient examples, 253 Santayana, George, 214 Schroeder, Jeanne, 229 scripts: as basis for conventional behavior, 192-93; and personal narratives, 200; and psychotherapy, 201-3. See also stock stories Searle, John, 295-96n.6 self-consciousness: critical, 135-37

Rousseau, Jean-Jacques, 2

self-criticism, 134–35, 142 self-mastery, 284 self-reference, 139; problem of, 125-29, 133; and self-criticism, 129-38 semiotic analysis. See conceptual oppositions semiotics. See rhetoric sexuality: discourses on, 281; Foucault's view of, 325-26n.29; problematization of, 283-84 sexual selection, 83-85 sign: deconstructive theory of, 302n.25 signals, 60 slogans, 81 Slotkin, Richard, 205 social memory, 203-4 sociology of knowledge, 132 software: as metaphor, xi-xii sõphrosynē, 284 spandrels, 72-73, 291, 299n.28 speciation, 91 Sperber, Dan, 44, 86, 303n.32 stereotypes, 252; gender, 224-26, 252; of homosexuals, 233-34; patriarchal, 228stock characters, 198-99 stock market, 84 stock stories, 194-98, 203-5, 213 stories: and spread of memes, 78, 81; transformations of, 54, 307n.24. See also stock stories structuralism, 320n.7; and conceptual oppositions, 217-19 subcategorization, 221-23 subindividual entities, x supraindividual entities, x, 10, 12, 174 symbionts: memes as, 62, 287 symbiosis, 62 symbol, 19, 275; as meme, 47. See also metaphors; metonymic models symbolic forms, 102, 104, 105, 111-12, 114, 186, 309n.2

Talmud, 1, 6 taste, articulation of, 28–29 Taylor, Charles, 264 technology, 6–7; benefits of, 283; changes in, 94; development of, 53, 282; influence on daily life of, 274–75; pejorative versus ambivalent, 282–83; as vehicle for memes, 54–56

teleology, 41 television, 80-81 Thanksgiving myth, 208 Thébaud, Jean-Loup, 154-55 Thomas, Clarence, 197 Thompson, John, 102, 104, 132, 310n.3; theory of ideology, 111-13 tools: language as, 24; as metaphor, 1-2; as necessary for making other tools, 1; of understanding, 1-2, 3, 23, 24-26, 38-39, 45, 96-97, 104, 108, 126-29, 139-40, 286 Torres, Gerald, 219, 224, 225 tort doctrines, 236 traditions: diachronic and synchronic views of, 49-50; dynamic nature of, 95-96; Gadamer's theory of, 7-10, 13, 14, 263, 297; institutional authority of, 85-86; Marx's view of, 182, 183 transcendence: of value versus of position, 157-61 transcendental arguments, 145-50 transcendental subject, 12, 13, 14 transcendent ideals, 163-67, 168 transcendent values, 144-50; and cultural relativism, 150-55; defined, 144-45; development of, 167-68; and historicism, 169-70; as indeterminate standard, 151; justice as, 143, 144-45; love as, 314n.1; and positive norms, 148, 155-57; and pragmatism, 167-69; solidarity as, 313-14n.1; truth as, 143, 144-45 transformation, problem of, 9-10 transmission, problem of, 9, 10 truth: adversarial view of, 246-47; and ideology, 126; Foucault's view of, 264, 280; transcendent ideal of, 146, 147, 149; as transcendent value, 143, 144-49 Turner, Stephen, 8, 9 Tversky, Amos, 175, 176, 183, 253, 258

unconscious: universal structures of, 12, 217–19 understanding. See cultural understanding The Uses of Pleasure (Foucault), 283–84

values: cultural tools used to express, 27–31. *See also* articulation of human values; transcendent values
Vico, Giambattista, 2, 6, 34–35, 243
viruses: memes as, 60–61, 287

war: as metaphor for rational argument, 245-47 Washington, George, 255 Weber, Max, 281, 282 Wilson, Edward O., 65, 67 wishful thinking, 176-77, 178, 310n.16 Wittgenstein, Ludwig, 33, 186, 250, women: associations, 34; stereotypes of, 224-26, 231-32

Wong, David B., 315n.6 writing: invention of, 78 written cultures, 79

yin and yang, 234

Zeitgeist, 95, 96; as effect of cultural software, 276 Zenger trial, 236-38 Zeno, 266