

---

# **Making It Explicit**

**Reasoning, Representing,  
and Discursive Commitment**

---

---

Robert B. Brandom

Harvard University Press  
Cambridge, Massachusetts  
London, England

---

# To Wilfrid Sellars and Richard Rorty

without whom most of it would not even be *implicit*

Copyright © 1994 by the President and Fellows of Harvard College  
All rights reserved  
Printed in the United States of America  
Page xxiv constitutes an extension of the copyright page.  
Fourth printing, 2001

First Harvard University Press paperback edition, 1998

*Library of Congress Cataloging-in-Publication Data*

Brandom, Robert.

Making it explicit : reasoning, representing, and discursive commitment /  
Robert B. Brandom.

p. cm.

Includes bibliographical references and index.

ISBN 0-674-54319-X (cloth)

ISBN 0-674-54330-0 (pbk.)

1. Language and languages—Philosophy. 2. Semantics (Philosophy)  
3. Pragmatics. 4. Representation (Philosophy) I. Title.

P106.B694 1994

121'.68—dc20

93-50631

---

. . . both a new world,  
and the old made explicit . . .

We shall not cease from exploration  
And the end of all our exploring  
Will be to arrive where we started  
And know the place for the first time.

T. S. ELIOT, "Four Quartets"



---

# Contents

---

*Preface* xi

## PART ONE

|  |           |
|--|-----------|
| <b>1. Toward a Normative Pragmatics</b>                          | <b>3</b>  |
| I. Introduction  | 3         |
| II. From Intentional State to Normative Status                   | 7         |
| III. From Norms Explicit in Rules to Norms Implicit in Practices | 18        |
| IV. From Normative Status to Normative Attitude                  | 30        |
| V. From Assessment to the Social Institution of Norms            | 46        |
| VI. From Intentional Interpretation to Original Intentionality   | 55        |
| <i>Appendix: Wittgenstein's Use of <i>Regel</i></i>              | 64        |
| <b>2. Toward an Inferential Semantics</b>                        | <b>67</b> |
| I. Content and Representation                                    | 67        |
| II. The Priority of the Propositional                            | 79        |
| III. Conceptual Classification and Inference                     | 85        |
| IV. Material Inference, Conceptual Content, and Expression       | 94        |
| V. Circumstances and Consequences of Application                 | 116       |
| VI. Conclusion   | 132       |

|   |            |
|---|------------|
| <b>3. Linguistic Practice and Discursive Commitment</b>   | <b>141</b> |
| I. Intentional States and Linguistic Practices  | 141        |
| II. Deontic Status and Deontic Attitudes  | 157        |
| III. Asserting and Inferring  | 167        |
| IV. Scorekeeping: Pragmatic Significance and Semantic Content   | 180        |
| <b>4. Perception and Action: The Conferral of Empirical and Practical Conceptual Content</b>                      | <b>199</b> |
| I. Assertions as Knowledge Claims   | 199        |
| II. Reliability   | 206        |
| III. Observation Reports and Noninferential Authority   | 213        |
| IV. Rational Agency   | 229        |
| V. Practical Reasoning: Inferences from Doxastic to Practical Commitments   | 243        |
| VI. Intentions  | 253        |
| PART TWO  |            |
| <b>5. The Expressive Role of Traditional Semantic Vocabulary: 'True' and 'Refers'</b>                             | <b>275</b> |
| I. From Inference to Truth, Reference, and Representation   | 275        |
| II. Truth in Classical Pragmatism   | 285        |
| III. From Pragmatism to Prosentences  | 299        |
| IV. Reference and Anaphorically Indirect Descriptions   | 305        |
| V. The Function of Traditional Semantic Vocabulary Is Expressive, Not Explanatory                                 | 322        |
| <b>6. Substitution: What Are Singular Terms, and Why Are There Any?</b>   | <b>334</b> |
| I. Multivalued Logic and Material Inference   | 334        |
| II. Substitution, Sentential Embedding, and Semantic Roles  | 346        |
| III. Subsential Expressions   | 360        |
| IV. What Are Singular Terms?  | 367        |
| V. Why Are There Singular Terms?  | 376        |
| VI. Objections and Replies  | 384        |
| VII. Conclusion   | 399        |
| <i>Appendix I: From Substitutional Derivation of Categories to Functional Derivation of Categories</i>            | 404        |
| <i>Appendix II: Sentence Use Conferring the Status of Singular Terms on Subsential Expressions—An Application</i> | 409        |
| <b>7. Anaphora: The Structure of Token Repeatables</b>  | <b>413</b> |
| I. Frege's <i>Grundlagen</i> Account of Picking Out Objects   | 413        |
| II. Definite Descriptions and Existential Commitments   | 432        |
| III. Substitution, Token Recurrence, and Anaphora   | 449        |
| IV. Deixis and Anaphora   | 459        |

|   |            |
|---|------------|
| V. Interpersonal Anaphora and Communication   | 473        |
| <i>Appendix: Other Kinds of Anaphora—Paychecks, Donkeys, and Quantificational Antecedents</i>   | 490        |
| <b>8. Ascribing Propositional Attitudes: The Social Route from Reasoning to Representing</b>  | <b>495</b> |
| I. Representation and <i>De Re</i> Ascription of Propositionally Contentful Commitments   | 495        |
| II. Interpretation, Communication, and <i>De Re</i> Ascriptions   | 508        |
| III. <i>De Re</i> Ascriptions and the Intentional Explanation of Action   | 520        |
| IV. From Implicit Attribution to Explicit Ascription  | 529        |
| V. Epistemically Strong <i>De Re</i> Attitudes: Indexicals, Quasi-Indexicals, and Proper Names  | 547        |
| VI. The Social-Perspectival Character of Conceptual Contents and the Objectivity of Conceptual Norms  | 584        |
| <i>Appendix: The Construction and Recursive Interpretation of Iterated Ascriptions That Mix <i>De Dicto</i> and <i>De Re</i> Content Specifications</i> | 608        |
| <b>9. Conclusion</b>  | <b>614</b> |
| I. Two Concepts of Concepts   | 614        |
| II. Norms and Practices   | 623        |
| III. We Have Met the Norms, and They Are Ours   | 639        |
| <i>Abbreviations</i>  | 653        |
| <i>Notes</i>  | 655        |
| <i>Index</i>  | 717        |





---

# Preface

---

We work in the dark—we do what we can—we give what we have. Our doubt is our passion and our passion is our task. The rest is the madness of art.

HENRY JAMES, "The Middle Years"

---

That old philosopher Fred Allen used to say he could not understand why someone would spend years writing a novel, when for a few dollars you could buy one practically anywhere. A similar remark might be made about contributions to that peculiar genre of creative nonfiction writing to which philosophical works such as this one belong. This book is an investigation into the nature of *language*: of the social practices that distinguish us as rational, indeed logical, concept-mongering creatures—knowers and agents. This is of course a topic that has been much explored by philosophers, both the mighty dead and the ablest contemporary thinkers. Surrounded as we are by the riches they have bequeathed, it is hard to avoid asking why one should bother reading—let alone writing—yet another such work. This question may seem all the more urgent inasmuch as it is acknowledged (indeed, some pains are taken to show) that the basic building blocks out of which this account is constructed—its motivating insights, commitments, and strategies—are not novel or original.

Still, though the ways of thinking and talking about thinking and talking presented here arise naturally out of a reading of the philosophical tradition (above all Kant, Frege, and Wittgenstein) and of its development by more recent thinkers, both that tradition and its significance for us today are by current standards seen decidedly on a slant. As a result, the story told in these pages comes at familiar things from an unfamiliar direction. Its promise lies

in the sort of added depth and dimension that only binocular vision affords; that is the point of laying a substantially different conceptual perspective alongside our more accustomed line of sight. In keeping with this understanding of the sort of payoff that can be hoped for, the body of the work aims to set criteria of adequacy for a theory of discursive practice, motivate the approach adopted, work the model out in detail, and apply it. The idea is to show what kind of understanding and explanatory power one gets from talking this way, rather than to argue that one is somehow rationally *obliged* to talk this way.

Of course I take it that the claims made in what follows are true; I endorse those assertions; they express my commitments. One of the central tenets of the account of linguistic practice put forward here is that the characteristic *authority* on which the role of assertions in communication depends is intelligible only against the background of a correlative *responsibility* to vindicate one's entitlement to the commitments such speech acts express. It is possible to secure entitlement to the commitments (assertional, inferential, and referential) implicit in an idiom without gainsaying the possibility of entitlement to a different one. But even such a modest justificatory project is of interest only to someone who both understands the commitments in question and has some reason to *want* to become entitled to talk in ways that presuppose them.

Both of these ends are served by starting the story with some historical lessons. Accordingly, Chapters 1 and 2 form an entrance hall to the rest of the edifice, one whose main architectural features are made more noticeable by the judicious placement of ancestor portraits. The same figure appears on many walls, almost always recognizable, but often portrayed from unusual vantage points (from behind, from above) or highlighting something other than the familiar face. In particular, the portrait of Frege will seem to some to be like one of those odd photographs of a reclining figure taken with the lens so close that the subject's left foot assumes gigantic proportions, dwarfing the rest of the individual, whose head and torso dwindle to the dimensions of insignificant appendages. Nonetheless, the tradition that is retrospectively constituted by the unusual emphases and filiations to be found here is meant to be coherent and compelling in its own terms. It is not just whatever rewriting of the history of philosophy happens to be needed to make the waning years of the twentieth century safe for the views I put forward. Rather, those views have the shape they do because of this reading of how we got to where we are.

One of the overarching methodological commitments that orients this project is to explain the *meanings* of linguistic expressions in terms of their *use*—an endorsement of one dimension of Wittgenstein's pragmatism. For although he drove home the importance of such an approach, other features of his thought—in particular his theoretical quietism—have discouraged his admirers from attempting to work out the details of a theory of meaning or,

for that matter, of use. One result has been a substantial disjunction between *semantic* theorizing (about the sorts of contents expressed by various locutions), on the one hand, and *pragmatic* theorizing (about the linguistic practices in which those locutions are employed), on the other. The explanatory strategy pursued here is to begin with an account of social practices, identify the particular structure they must exhibit in order to qualify as specifically *linguistic* practices, and then consider what different sorts of semantic contents those practices can confer on states, performances, and expressions caught up in them in suitable ways. The result is a new kind of conceptual-role semantics. It is at once firmly rooted in actual practices of producing and consuming speech acts, and sufficiently finely articulated to make clear how those practices are capable of conferring the rich variety of kinds of content that philosophers of language have revealed and revealed in.

Claims about the relations between meaning and use have a clear sense only in the context of a specification of the vocabulary in which that use is described or ascribed. At one extreme, use clearly determines meaning in the strongest possible sense if admissible specifications of the use can include such phrases as "using the word 'not' to express negation" or "using the term 'Julius Caesar' to refer to Julius Caesar." At an opposite extreme, if admissible specifications of use are restricted to descriptions of the movements of particles expressed in the vocabulary of physics, not only will the use, so described, fail to settle *what* is meant or expressed by various noises or inscriptions, it will fail to settle even *that* anything is meant or expressed by them. The specification of use employed here is neither so generous as to permit semantic or intentional vocabulary nor so parsimonious as to insist on purely naturalistic vocabulary.

Instead, it makes essential use of *normative* vocabulary. The practices that confer propositional and other sorts of conceptual content implicitly contain norms concerning how it is *correct* to use expressions, under what circumstances it is *appropriate* to perform various speech acts, and what the *appropriate* consequences of such performances are. Chapter 1 introduces and motivates this normative pragmatics, which is rooted in considerations advanced by Kant, Frege, and Wittgenstein. No attempt is made to eliminate, in favor of nonnormative or naturalistic vocabulary, the normative vocabulary employed in specifying the practices that are the use of a language. Interpreting states, performances, and expressions as semantically or intentionally contentful is understood as attributing to their occurrence an ineliminably normative pragmatic significance.

Though this normative dimension of linguistic practice is taken to be ineliminable, it is not treated as primitive or inexplicable. It is rendered less mysterious in two ways. First, linguistic norms are understood as *instituted* by social-practical activity. The pragmatic significances of different sorts of speech acts are rendered theoretically in terms of how those performances

affect the *commitments* (and *entitlements* to those commitments) acknowledged or otherwise acquired by those whose performances they are. The norms implicit in linguistic practice are accordingly presented in a specifically *deontic* form. But these deontic statuses are understood in turn as a form of *social* status, instituted by the practical attitudes of those who attribute and acknowledge such statuses.

The natural world does not come with commitments and entitlements in it; they are products of human activity. In particular, they are creatures of the *attitudes* of taking, treating, or responding to someone in practice as committed or entitled (for instance, to various further performances). Mastering this sort of norm-instituting social practice is a kind of practical know-how—a matter of keeping deontic *score* by keeping track of one's own and others' commitments and entitlements to those commitments, and altering that score in systematic ways based on the performances each practitioner produces. The norms that govern the use of linguistic expressions are implicit in these deontic scorekeeping practices.

The second way norms are rendered less mysterious is by explaining exactly what is expressed by normative vocabulary. Beginning with basic deontic scorekeeping attitudes and the practices that govern them, an account is offered of how locutions must be *used* in order to express explicitly the very normative notions—*is committed*, *is permitted*, *ought*, and so on—that are appealed to in laying out the normative pragmatics. This is an explication of explicitly normative conceptual contents in terms of implicitly normative practices, rather than a reduction of normative terms to nonnormative ones. It illuminates the normative dimension of discursive practice in line with the methodological principle that implicit structures are often best understood by looking at how they can be made explicit.

The first step in the project is accordingly the elaboration of a pragmatics (a theory of the use of language) that is couched in terms of practical scorekeeping *attitudes* of attributing and acknowledging deontic *statuses* of commitment and entitlement. The pragmatic significance of performances—eventually, speech acts such as assertions—is then understood to consist in the difference those performances make to the commitments and entitlements attributed by various scorekeepers. The next step is to say what structure such a set of social practices must have in order to qualify as specifically *discursive* practice. This is a matter of moving from pragmatics to semantics. The defining characteristic of discursive practice is the production and consumption of specifically *propositional* contents. It is argued in Chapter 2 that propositional contentfulness should be understood in terms of *inferential* articulation; propositions are what can serve as premises and conclusions of inferences, that is, can serve as and stand in need of *reasons*. Chapter 3 describes (in deontic scorekeeping terms) a model of social practices of giving and asking for reasons—specifically *linguistic* discursive prac-

tices, which suffice to confer propositional contents on states, attitudes, performances, and expressions that play suitable roles in those practices.

This account of the conferral of semantic content by inferentially articulated, social scorekeeping practice is further generalized in two different directions. First, it is shown how this model can be applied in order to understand not only linguistic meanings but intentional contents generally. The propositional contentfulness of *beliefs*, no less than of *claims*, should be understood in terms of their role in reasoning of various kinds. The inferentially articulated commitments expressed by assertional speech acts are *doxastic* commitments. Much of the theoretical work done by the concept of belief can be done instead by appeal to this sort of deontic status, and to the practical scorekeeping attitudes of acknowledging or undertaking such commitments. A social, linguistic account of *intentionality* is accordingly elaborated in Chapter 3.

It is extended in Chapter 4 to incorporate treatments of *perception* and *action* and of the contribution those phenomena make to the empirical and practical dimensions of the propositional contents of the states, acts, and attitudes involved in them. It is not denied that it makes sense to talk about nonlinguistic creatures as having intentional states, but it is claimed that our understanding of such talk is parasitic on our understanding of the sort of full-blooded linguistic intentionality characteristic of states and attitudes that only beings who engage in discursive social practices can have. This story amounts, then, to an account of the relations of mindedness—in the sense of sapience rather than mere sentience—to behavior. As in the parallel case of meaning and use, the clarification of these relations must begin with a determination of what vocabulary it is admissible to use in specifying the relevant behavior; again there is a spectrum of possibilities, from allowing intentional vocabulary with semantic locutions, such as “acting as if one believed that snow is white,” ranging down to restrictions to physicalistic or other naturalistic vocabulary, such as “one’s left wrist rotating twenty degrees.” The *via media* pursued here eschews intentional or semantic specifications of behavior but permits normative and therefore social specifications of what is in fact linguistic behavior.

Where the first sort of generalization involves moving from consideration of *language* to consideration of *mind*, from talking to thinking and believing, the second involves moving from an account of the practices that constitute treating something as *propositionally* contentful to the practices that constitute treating something as *conceptually* contentful in a broader sense. In Chapter 6 the notion of *substitution* and substitutional inferences is used to show how expressions such as singular terms and predicates, which cannot directly play the inferential role of premise or conclusion in an argument, nonetheless can play an *indirectly* inferential role in virtue of their systematic contributions to the directly inferential roles of sentences in which they

occur. In Chapter 7 the notion of *anaphora* (whose paradigm is the relation between a pronoun and its antecedent) and anaphoric inheritance of substitutional commitment is used to show how even unrepeatable expressions such as demonstrative tokenings play substitution-inferential roles and hence express conceptual contents. The result is a kind of conceptual-role semantics that is distinguished first by the nature of the functional system with respect to which such roles are individuated and attributed: what is appealed to is role in the implicitly normative linguistic social practices of a community, rather than the behavioral economy of a single individual. It is also different from familiar ways of using the notion of conceptual role in conceiving of the conceptual in terms of specifically *inferential* articulation, and in its elaboration of the fundamental *substitutional* and *anaphoric* substructures of that inferential articulation.

This semantic explanatory strategy, which takes *inference* as its basic concept, contrasts with the one that has been dominant since the Enlightenment, which takes *representation* as its basic concept. The inferentialist approach is by no means without precedent—though it has been largely a minority platform. Indeed, the distinction canonically drawn between Continental rationalists such as Spinoza and Leibniz, on the one hand, and British empiricists such as Locke and Hume, is for many purposes more perspicuously rendered as a distinction between those endorsing an inferentialist order of explanation and those endorsing a representationalist order of explanation. The elements of the contemporary inferentialist program are extracted (in Chapter 2) from Frege of the *Begriffsschrift*, Sellars, and some of Dummett's writings.

The complementary theoretical semantic strategies of representationalism and inferentialism are bound by the same pair of general explanatory obligations: to explicate the concept treated as primitive, and to offer an account of other semantic concepts in terms of that primitive. The representationalist tradition has developed good answers to the second sort of concern, primarily by employing a variety of set-theoretic methods to show how proprieties of inference can be determined by representational properties of the claims that serve as their premises and conclusions. The explanatory challenge to that tradition lies rather in the first sort of demand, in saying what it is for something to have representational content, and in what the grasp or uptake of that content by speakers and thinkers consists. As the inferentialist program is pursued here, the proprieties of inference that serve as semantic primitives are explicated in the pragmatics; they are implicit in the practices of giving and asking for reasons. The major explanatory challenge for inferentialists is rather to explain the representational dimension of semantic content—to construe *referential* relations in terms of *inferential* ones.

The second part of the book responds to this challenge. Chapter 5 explains the expressive role of traditional representational semantic vocabulary. An account is offered there of the use of the sort of expression of which 'true'

and 'refers' are paradigmatic. Following the lead of the prosentential approach to truth, the key semantic concept employed in that unified account is *anaphora*. Chapter 7 then explains anaphoric relations in terms of the substitution-inferential structure of discursive scorekeeping elaborated in Chapter 6. Chapter 6 also offers an account in those terms of what it is for claims—which are understood in the first instance (in Chapter 3) as what can serve as premises and conclusions of inferences—to be and be understood to be about *objects*, and to characterize them as having properties and standing in relations.

The primary treatment of the representational dimension of conceptual content is reserved for Chapter 8, however. There the representational properties of semantic contents are explained as consequences of the essentially *social* character of inferential practice. Words such as the 'of' that expresses intentional directedness, and 'about' and 'represents' in their philosophically significant uses, have the expressive role they do—making representational relations explicit—in virtue of the way they figure in *de re* ascriptions of propositional attitudes. These are the tropes used to say explicitly what someone is thinking *about*, what a belief *represents*, what a claim is true *of*. Chapter 8 offers a discursive scorekeeping account of the practices that constitute using locutions to express such *de re* ascriptions, and hence of how expressions must be used in order to mean 'of', 'about', or 'represents'. This account of what is expressed by the fundamental explicitly representational locutions makes possible an explanation of the *objectivity* of concepts. It takes the form of a specification of the particular sort of inferential structure social scorekeeping practices must have in order to institute objective norms, according to which the correctness of an application of a concept answers to the facts about the object to which it is applied, in such a way that anyone (indeed *everyone*) in the linguistic community may be wrong about it.

In summary, in the theoretical place usually occupied by the notion of *intentional states*, the pragmatics presented here elaborates a conception of *normative statuses*; in the place usually occupied by the notion of *intentional interpretation*, it puts *deontic scorekeeping*—that is, the social practices of attributing and acknowledging commitments and entitlements, which implicitly institute those statuses. The theoretical work typically done by semantic assessments according to correctness of *representation* and satisfaction of *truth* conditions is done by assessments of proprieties of *inference*. Semantic articulation is attributed and acknowledged by keeping score not only of directly inferential commitments, which relate sentential (that is, claimable or believable) contents, but also of indirectly inferential *substitutional* and *anaphoric* commitments, which relate the subsentential contents of expressions of other grammatical categories.

The pragmatics and semantics maintain particularly intimate relations throughout. The aim is always to show how some bit of vocabulary must be *used*—the significance its utterance must have in various circumstances, the

practical scorekeeping attitudes its usage must elicit and be elicited by—in order for it to express a certain kind of semantic content: to be being taken or treated in practice by the linguistic community *as* a conditional, a singular term, a bit of normative vocabulary, a propositional-attitude-ascribing locution, and so on. A fundamental methodological criterion of adequacy of the account is that the theorist not attach semantic contents to expressions by *stipulation*; it must always be shown how such contents can be *conferred* on expressions by the scorekeeping activities the theorist attributes to the linguistic practitioners themselves. That is, the aim is to present conditions on an interpretation of a community as discursive scorekeepers that are sufficient (though perhaps not necessary) to ensure that interpreting the community as engaged in *those* implicitly normative practices *is* interpreting them as taking or treating their speech acts as expressing the sorts of semantic contents in question.

The obligation to say what it is about the use of locutions in virtue of which they express various sorts of content dictates that the master concept articulating the relation between the pragmatic and semantic portions of the theory is that of *expression*. To express something is to make it *explicit*. What is explicit in the fundamental sense has a *propositional* content—the content of a claim, judgment, or belief (claimable, judgeable, believable contents). That is, making something explicit is *saying* it: putting it into a form in which it can be given as a reason, and reasons demanded for it. Putting something forward in the explicit form of a claim is the basic move in the game of giving and asking for reasons.

The relation of expression between what is *implicit* in what practitioners *do* and what is *explicit* in what they *say* structures the story told here at two different levels. At the basic level, the question is how the capacity to entertain principles, and so to know *that* something is the case, arises out of the capacity to engage in practices—to know *how* to do something in the sense of being *able* to do it. What must practitioners be able to *do* in order to be able thereby to *say that* things are thus and so—that is, to express something explicitly? The explanatory force of a response to this question can be judged by the constraints that are acknowledged on the vocabulary in which those practical capacities are specified; normative vocabulary is employed here, but intentional vocabulary (which would permit at the outset the ascription of propositionally contentful states, attitudes, and performances) is not. The first level of the account of expression accordingly consists in explaining—making theoretically explicit—the implicit structure of linguistic practices in virtue of which they count as making anything explicit at all.

The second level of the account of expression consists in working out a theory of the expressive role distinctive of *logical* vocabulary. The claim is that logical vocabulary is distinguished by its function of expressing explicitly *within* a language the features of the use of that language that confer



conceptual contents on the states, attitudes, performances, and expressions whose significances are governed by those practices. Conditionals serve as a paradigm illustrating this expressive role. According to the inferential approach to semantics and the deontic scorekeeping approach to pragmatics, practitioners confer determinate propositional contents on states and expressions in part by their scorekeeping practice of treating the acknowledgment of one doxastic commitment (typically through assertional utterance of a sentence) as having the pragmatic significance of an undertaking of further commitments that are related to the original commitment as its inferential consequences. At the basic level, treating the claim expressed by one sentence as an inferential consequence of the claim expressed by another sentence is something practitioners can *do*, and it is because such practical attitudes can be implicit in the way they respond to each other's performances that their sentences come to mean what they do. With the introduction of conditional locutions linking sentences, however, comes the expressive power to *say* explicitly *that* one claim is a consequence of another. The expressive role distinctive of conditionals is making *implicit* inferential commitments *explicit* in the form of declarative sentences, the assertion of which acknowledges a propositionally contentful doxastic commitment. In a similar way, at the basic level, scorekeepers can treat the claims expressed by two sentences as incompatible—namely by treating commitment to one as in practice precluding entitlement to the other. The introduction of a locution with the expressive power of negation makes it possible to express such implicit practical scorekeeping attitudes explicitly—by *saying* that two claims are incompatible (one entails the negation of the other). Identity and quantificational expressions are analyzed on this model as making explicit the substitutional relations characteristic of singular terms and predicates respectively, and further locutions are considered that play a corresponding expressive role in making anaphoric relations explicit.

So an *expressive* theory of logic is presented here. On this view, the philosophical significance of logic is not that it enables those who master the use of logical locutions to *prove* a special class of claims—that is, to entitle themselves to a class of commitments in a formally privileged fashion. The significance of logical vocabulary lies rather in what it lets those who master it *say*—the special class of claims it enables them to express. Logical vocabulary endows practitioners with the expressive power to make explicit as the contents of claims just those implicit features of linguistic practice that confer semantic contents on their utterances in the first place. Logic is the organ of semantic self-consciousness. It brings out into the light of day the practical attitudes that determine the conceptual contents members of a linguistic community are able to express—putting them in the form of explicit claims, which can be debated, for which reasons can be given and alternatives proposed and assessed. The formation of concepts—by means of which practitioners can come to be aware of anything at all—comes itself to

be something of which those who can deploy logical vocabulary can be aware. Since plans can be addressed to, and intentional practical influence exercised over, just those features of things of which agents can become explicitly aware by the application of concepts, the formation of concepts itself becomes in this way for the first time an object of conscious deliberation and control.

Explaining the features of the use of *logical* vocabulary that confer its characteristic sort of semantic content is accordingly explaining how the sort of expressive power the theorist requires to explain the features of the use of *nonlogical* vocabulary that confer semantic content on it can become available to those whose linguistic practice is being theorized about. It is this fact that sets the expressive scope of the project pursued here. The aim is twofold: to make explicit deontic scorekeeping social practices that suffice to confer conceptual contents on nonlogical sentences, singular terms, and predicates in general; and to make explicit the deontic scorekeeping social practices in virtue of which vocabulary can be introduced as playing the expressive roles characteristic of a variety of particular logical locutions. How much logical vocabulary is worth reconstructing in this fashion? In this project, neither more nor less than is required to make explicit *within* the language the deontic scorekeeping social practices that suffice to confer conceptual contents on nonlogical vocabulary in general. At that point it will have been specified what practices a theorist must attribute to a community in order to be interpreting its members as engaging not just in specifically *linguistic* practices but in linguistic practices that endow them with sufficient expressive power to *say* how their practices confer conceptual content on their states, attitudes, performances, and expressions. That is, they will be able to express the theory offered here.

To make the semantic theory explicit requires logical vocabulary capable of expressing inferential, substitutional, and anaphoric relations. This vocabulary corresponds pretty well to the language of standard first-order logic, with the addition of classical semantic vocabulary. To make the pragmatic theory explicit requires logical vocabulary expressing the endorsement of norms generally, and the attribution and acknowledgment of the deontic statuses of commitment and entitlement in particular. The discussion of action in Chapter 4 includes an account of the use of vocabulary that makes norms explicit, and Chapter 8 explains how the social-practical attitudes of attributing and acknowledging deontic statuses (paradigmatically doxastic commitment) are made explicit by the use of propositional-attitude-ascribing locutions such as the regimented '*. . . is committed to the claim that . . .*,' which does duty here for '*. . . believes that . . .*'. Along this expressive dimension, the project eats its own tail, or lifts itself up by its own bootstraps—presenting an explanation of what it is to say something that is powerful enough to explain what it itself is saying.

Interpreting the members of a community as engaging in specifically

*discursive* practices, according to the view put forward here, is interpreting them as engaging in social practices that include treating some performances as having the pragmatic significance of *assertions*. For it is in terms of the constellation of inferentially articulated commitments and entitlements characteristic of the making (staking) of claims that the notion of specifically *propositional* contentfulness is to be understood. Since all other varieties of conceptual contentfulness derive (substitutionally) from the propositional, this is to say that the application of concepts is a linguistic affair—not in the sense that one must be talking in order to do it, but in the sense that one must be a player of the essentially linguistic game of giving and asking for reasons in order to be able to do it. There can be sets of practices that are linguistic in this sense but that do not incorporate the expressive resources provided by logical vocabulary. Indeed, the way the use of such vocabulary in making explicit what is implicit in the use of nonlogical vocabulary is specified is by showing what would be required to introduce vocabulary with that expressive function into idioms that did not already contain it. The contribution made by logical locutions to the reflective processes in virtue of which the evolution of our concepts and commitments qualifies as *rational* is so important, however, that linguistic practices that at least permit their introduction form a special class. In a weak sense, any being that engages in linguistic practices, and hence applies concepts, is a *rational* being; in the strong sense, rational beings are not only linguistic beings but, at least potentially, also *logical* beings. This is how we should understand ourselves: as beings that meet this dual expressive condition.

It turns out that there is a surprising connection between being a rational creature—in the sense that includes the possibility of using the expressive resources of specifically *logical* vocabulary to reflect on one's conceptual content-conferring linguistic practices—on the one hand, and the structure of the *facts* that make up the world one can become aware of by applying those concepts, on the other. Rational beings live in a world of propertied and related particulars. Chapter 6 presents an expressive deduction of the necessity of this structure; it shows not why there is something rather than nothing but why what there is must come in the form of *things*; it shows why judgments or beliefs—the commitments expressed by claims—must in the basic case be about particulars (paradigmatically objects) and their properties and relations.

Particular objects are what is referred to by singular terms, and the demonstration proceeds by showing that the only semantically significant sub-sentential structure that is compatible with the introduction of logical vocabulary is one that decomposes basic sentences into singular terms and predicates. This would not be a surprising result if the logical vocabulary appealed to included identity and quantificational locutions, for (it will be argued) these have precisely the expressive role of making explicit in the form of claims the substitution-inferential commitments characteristic of

singular terms and predicates. But the result presented here is much stronger: any discursive practices that permit the introduction even of *sentential* logical operators such as negation and conditionals require that any *subsential* substitutional structure be of the term-predicate variety. Thus the investigation of the nature and limits of the explicit expression in principles of what is implicit in discursive practices yields a powerful transcendental argument—a formal answer to the question, Why are there objects? that turns on a deep relation between the expressive capacities required to think critically about the inferential connections among claims and the structures in virtue of which those claims are properly understood as characterizing objects as having properties and standing in relations.

This is a long book. Its length is a consequence of the demands made by its governing methodological aspirations: to eschew representational primitives, to show how content is related to use, and to achieve self-referential expressive completeness. The first is pursued by elaborating inferentialist and expressivist alternatives to the representationalist idiom for thinking and talking about thinking and talking that has been so well worked out over the last three centuries. The aim is not to replace that familiar idiom but to enrich it. The promised enrichment is of two sorts. First, there is the greater depth of field afforded by the stereoscopic vision made available by an alternative to familiar ways of talking about intentional phenomena. Second, there is the grounding and illumination of representational tropes secured by displaying the implicit features of discursive practice that are expressed explicitly by their use. Doing this requires that both the pragmatics and the semantics be developed in a reasonable amount of detail. The account of norm-instituting social practices must appeal to capacities that are plausibly available in primitive prelinguistic cases, and yet provide raw materials adequate for the specification of sophisticated linguistic practices, including logical ones. The account of the semantic contents conferred by those practices must encompass expressions of grammatical categories that are reasonably well understood already within the representationalist tradition—for example predicates, definite descriptions, proper names, familiar sorts of logical expressions, and whatever other kinds of locutions are required to make the processes by which content is conferred explicit within the linguistic practices being modeled.

Chapters 3 and 4 present the core theory—the model according to which a pragmatics specifying the social practices in which conceptual norms are implicit and a broadly inferential semantics are combined. It is here that sufficient conditions are put forward for the practices a community is interpreted as engaging in to count as according performances the pragmatic significance characteristic of *assertions*—and hence for those practices to count as conferring specifically *propositional* contents. Everything else in the book either leads up to the presentation of this model or elaborates and

extracts consequences from it. These chapters can be read on their own; the cost of omitting the first two chapters, which motivate the approach to pragmatics and semantics pursued here (in part by a rational reconstruction of the history of discussions of conceptual norms and contents), is that without this conceptual and historical background, one will not understand why things are done as they are here, rather than in some more familiar way. The time spent developing that motivation, however, means that one must wait a while for the actual theory to show itself.

The cost of missing Part 2 (Chapters 5 through 8) would be largely that one would then not see what the model can do, what it is good for. The most essential bit is Chapter 8, for that is where the *representational* dimension of discursive practice is explained in terms of the interaction of the *social* and the *inferential* articulation of the communication of reasons for belief and action. It is this interaction that is appealed to there also to make intelligible how *objective* norms come to apply to the essentially social statuses—paradigmatically the doxastic and practical propositionally contentful commitments that correspond to beliefs and intentions—and so underwrite such fundamental practices as assessing the *truth* of beliefs and the *success* of actions. The next most important part of the second half of the book is Chapter 6: the substitutional analysis is crucial to understanding how the inferential approach can generalize beyond sentences. And it would be truly a shame to miss the transcendental expressive argument for the existence of objects—the argument that (and why) the *only* form the world we talk and think of can take is that of a world of facts about particular objects and their properties and relations. It is worth keeping in mind Diderot's thought that one "must have gone deep into art or science to master their elements . . . The darkness of the beginnings lights up only toward the middle or the end."

The aim throughout is to present a unified vision of language and mind—one that starts with a relatively clear philosophical rationale and works it out in convincing detail, addressing a sufficiently wide range of potentially puzzling phenomena to engender confidence in its adaptability and power. It is animated by the ideal of the systematic philosophers of old: the invigorating clarifying prospect achievable by laying alongside our ordinary ways of talking and thinking an alternate idiom in which everything can be said. I am sensible, of course, of many ways in which this product falls short of that ideal. Particularly in matters of detail (but by no means there alone), a myriad of choices have had to be made at the cost of spurning attractive, perhaps ultimately superior, alternatives. The approach seldom dictates just one way of doing things. Yet the choice of which large limb to follow off the trunk of the tradition must be made on the basis of the tempting fruit to be seen on the smaller branches it supports. It can only be hoped that where upon closer inspection some of them are found wanting, the fundamental soundness of the tree is not impugned, but only the judgment of the gardener, who pruned

the better and nurtured the worse. As Johnson says in the Preface to his *Dictionary*: "A large work is difficult because it is large, even though all its parts might singly be performed with facility; where there are many things to be done, each must be allowed its share of time and labor in the proportion only which it bears to the whole; nor can it be expected that the stones which form the dome of the temple should be squared and polished like the diamond of a ring." Truth may or may not be in the whole, but understanding surely is.

I am grateful to the copyright holders for permission to reprint the following excerpts:

From Donald Davidson, *Actions and Events* (Oxford: Oxford University Press, 1980), by permission of Oxford University Press.

From Michael Dummett, *Frege: Philosophy of Language* (Cambridge, Mass.: Harvard University Press, copyright © 1991 by Michael Dummett), reprinted by permission of Harvard University Press and Gerald Duckworth & Co. Ltd.

From T. S. Eliot, "Burnt Norton" and "Little Gidding" in *Four Quartets*, copyright © 1943 by T. S. Eliot and renewed 1971 by Esme Valerie Eliot, reprinted by permission of Harcourt Brace & Company and Faber and Faber Ltd.

From Gottlob Frege, *The Foundations of Arithmetic*, trans. J. L. Austin (Evanston, Ill.: Northwestern University Press, 1978), by permission of the publisher.

From H. Hermes, F. Kambartel, and F. Kaulbach, eds., *Gottlob Frege: Posthumous Writings* (Chicago: University of Chicago Press, 1979), by permission of the University of Chicago Press and Blackwell Publishers.

From Wilfrid Sellars, "Inference and Meaning," *Mind* 62, no. 247 (1953), by permission of Oxford University Press.

From Wallace Stevens, "On the Road Home," *Collected Poems*, copyright 1942 by Wallace Stevens and renewed 1970 by Holly Stevens. Reprinted by permission of Alfred A. Knopf, Inc., and Faber and Faber Ltd.

Some of the material presented here has appeared elsewhere, in article form. In particular, "Inference, Expression, and Induction" (*Philosophical Studies* 54 [1988]: 257–285, reprinted by permission of Kluwer Academic Publishers) is essentially incorporated as Sections III–V of Chapter 2; "Asserting" (*Nous* 17, no. 4 [November 1983]: 637–650) is an early ancestor of the core of Chapter 3; and "Pragmatism, Phenomenalism, and Truth Talk" (*Midwest Studies in Philosophy* 12, "Realism and Antirealism" [1988]: 75–93) and "Reference Explained Away" (*The Journal of Philosophy* 81, no. 9 [September 1984]: 469–492) coincide extensively with the middle sections of Chapter 5.

During the long gestation of this work I have been helped immensely by discussions of these ideas with many other philosophers, who gave generously of their time and acumen. The confusions and unclaritys they have talked me through and out of are legion; those that remain belong on my account alone. It is impossible properly to acknowledge all these contributions, but some require particular mention. To begin with, I want to thank those with and to whom I most proximally say 'we': my colleagues in the Philosophy Department at the University of Pittsburgh, especially John McDowell, John Haugeland, and Nuel Belnap. The example, encouragement, and indulgence of this nearly ideal intellectual community made this book possible. I want also to express my appreciation for the exceptionally talented group of graduate students at Pitt who have shared this intellectual adventure over the years, particularly Mark Lance, Marc Lange, Danielle Macbeth, and Katarzyna Paprzycka. Finally, I am immeasurably grateful to Allan Gibbard for the gargantuan efforts he expended in mastering the ideas presented here, and for his suggestions as to how the final product might be improved.





---

# **PART ONE**

---



## Toward a Normative Pragmatics

---

An ounce of practice is worth a pound of precept.

ENGLISH PROVERB

---

### I. INTRODUCTION

#### 1. *Saying 'We'*

'We' is said in many ways. We may be thee and me. We may be all that talks or all that moves, all that minds or all that matters. Since these boundaries are elastic, we have a task of demarcation: telling who or what we are, distinguishing ourselves from the other sorts of objects or organisms we find in our world. Saying who we are can contract to an empty exercise in self-congratulation—a ritual rehearsal of the endless, pitiable disabilities of clockworks, carrots, cows, and the clan across the river. Such a mean-spirited version of the demarcational enterprise is not forced on us by the way things are, however.

For what we are is made as much as found, decided as well as discovered. The sort of thing we are depends, in part, on what we take ourselves to be. One characteristic way we develop and make ourselves into what we are is by expressing, exploring, and clarifying our understanding of what we are. Arbitrary distinctions of biology, geography, culture, or preference can be and have been seized on to enforce and make intelligible the crucial distinction between us and them (or it). But philosophical thought is coeval with the impulse to understand ourselves according to a more principled, less parochial story—and so to be a more principled, less parochial sort of being.

The wider perspective enjoined by principle poses the question, Who are we? in the form: What would have to be true—not only of the quaint folk across the river, but of chimpanzees, dolphins, gaseous extraterrestrials, or digital computers (things in many ways quite different from the rest of us)—for them nonetheless to be correctly counted among us? Putting the issue this way acknowledges an expansive demarcational commitment to avoid, as far as possible, requiring the sharing of adventitious stigmata of origin or material constitution. In understanding ourselves we should look to conditions at once more abstract and more practical, which concern what we are able to *do*, rather than where we come from or what we are made of. Candidates for recognition as belonging among us should be required to share only the fundamental abilities that make possible participation in those central activities by which we (thereby) define ourselves. How should we think of these?

The most cosmopolitan approach begins with a pluralistic insight. When we ask, Who are we? or What sort of thing are we? the answers can vary without competing. Each one defines a different way of saying 'we'; each kind of 'we'-saying defines a different community, and we find ourselves in many communities. This thought suggests that we think of ourselves in broadest terms as the ones who say 'we'. It points to the one great Community comprising members of all particular communities—the Community of those who say 'we' with and to someone, whether the members of those different particular communities recognize each other or not.

The reflexive character of the proposal that we use self-demarcation as the criterion by which we demarcate ourselves does not suffice to render it purely formal, however. It does not save us the trouble of contentful self-understanding. For until it has been specified in other terms what one must be able to do in order to count as "saying 'we'," demarcation by appeal to such attitudes remains an aspiration tacked to a slogan—empty, waiting for us to fill it. 'We'-saying of the sort that might be of demarcational interest is not a matter merely of the production of certain vocables—indeed perhaps the relevant kind of attitude is not a linguistic matter at all. Nor again does it consist simply in the engendering of warm mammalian fellow-feeling. Making explicit to ourselves who we are requires a theoretical account of what it is in practice to treat another as one of us.

## 2. *Sapience*

What is it we do that is so special? The answer to be explored here—a traditional one, to be sure—is that we are distinguished by capacities that are broadly cognitive. Our transactions with other things, and with each other, in a special and characteristic sense *mean* something to us, they have a *conceptual content* for us, we *understand* them in one way rather than another. It is this demarcational strategy that underlies the classical iden-

tification of us as *reasonable* beings. Reason is as nothing to the beasts of the field. We are the ones on whom reasons are binding, who are subject to the peculiar force of the better reason.

This force is a species of *normative* force, a rational 'ought'. Being rational is being bound or constrained by these norms, being subject to the authority of reasons. Saying 'we' in this sense is placing ourselves and each other in the space of reasons, by giving and asking for reasons for our attitudes and performances. Adopting this sort of practical stance is taking or treating ourselves as subjects of cognition and action; for attitudes we adopt in response to environing stimuli count as *beliefs* just insofar as they can serve as and stand in need of reasons, and the acts we perform count as *actions* just insofar as it is proper to offer and inquire after reasons for them. Our attitudes and acts exhibit an intelligible content, a content that can be grasped or understood, by being caught up in a web of reasons, by being inferentially articulated. Understanding in this favored sense is a grasp of reasons, mastery of proprieties of theoretical and practical *inference*. To identify ourselves as *rational*—as the ones who live and move and have our being in the space of reasons, and so to whom things can be intelligible—is to seize demarcationally on a capacity that might well be shared by beings quite different from us in provenance and demeanor.

Picking us out by our capacity for reason and understanding expresses a commitment to take *sapience*, rather than *sentience* as the constellation of characteristics that distinguishes us. Sentience is what we share with non-verbal animals such as cats—the capacity to be *aware* in the sense of being *awake*. Sentience, which so far as our understanding yet reaches is an exclusively biological phenomenon, is in turn to be distinguished from the mere reliable differential responsiveness we sentient beings share with artifacts such as thermostats and land mines. Sapience concerns understanding or intelligence, rather than irritability or arousal. One is treating something as sapient insofar as one explains its behavior by attributing to it intentional states such as belief and desire as constituting reasons for that behavior.

Another familiar route to understanding the sort of sapience being considered here for demarcational duty goes through the concept of *truth*, rather than that of *inference*. We are believers, and believing is taking-true. We are agents, and acting is making-true. To be sapient is to have states such as belief, desire, and intention, which are contentful in the sense that the question can appropriately be raised under what circumstances what is believed, desired, or intended would be *true*. Understanding such a content is grasping the conditions necessary and sufficient for its truth.

These two ways of conceiving sapience, in terms of inference and in terms of truth, have as their common explanatory target contents distinguished by their *propositional* form. What we can offer as a reason, what we can take or make true, has a propositional content: a content of the sort that we express by the use of declarative sentences and ascribe by the use of 'that' clauses.

Propositional contents stand in inferential relations, and they have truth conditions. One of the tasks of this work is to explain what it is to grasp specifically propositional contents, and so to explain who we are as rational or sapient beings. A central subsidiary task is accordingly to offer an account of the relation between the concepts of inference and truth, which complement one another and in some measure compete with one another for explanatory priority in addressing the issue of propositional contentfulness, and so of rationality.

### 3. *Intentionality*

The general self-understanding in view so far identifies us by our broadly cognitive capacities: We are makers and takers of reasons, seekers and speakers of truth. The propositional focus of the approach marks this understanding of intelligible contents as *discursive*. This conception, hal- lowed by ancient tradition, was challenged during the Enlightenment by a rival approach to cognitive contentfulness that centers on the concept of *representation*. Descartes's seminal demarcational story distinguishes us as representers—producers and consumers of representings—from a world of merely represented and representable things. The states and acts characteristic of us are in a special sense *of, about, or directed at* things. They are representings, which is to say that they have representative content. To have such a content is to be liable to assessments of *correctness* of representation, which is a special way of being *answerable or responsible* to what is represented.

Another task of this work is accordingly to address the question, How should the relation between representation—the master concept of Enlightenment epistemology—and the discursive concepts of reasons and truth be understood? One of the great strengths representationalist explanatory strategies have developed is the capacity to offer accounts of truth and goodness of inference. There are familiar set-theoretic routes that set out from representational primitives corresponding to subsentential linguistic expressions such as singular terms and predicates, lead to assignments of truth conditions to sentences compounded out of those expressions, and pass from there to determinations of which inferences are correct. While doubts have been raised, perhaps legitimately, about nearly every phase of this construction, no other semantic approach has been worked out so well.

Yet for all that, the primitives involved have never been well understood. Descartes notoriously fails to offer an account either of the nature of representational contents—of what the representingness of representings consists in—or of what it is to grasp or understand such contents, that is to say, of their intelligibility to the representer. He does not tell us what makes a rabbit-idea an idea of (or purporting to be of) rabbits, or of anything at all, nor what it is for the one whose idea it is to understand or take it *as* being of or about something. That things could be represented by and to the mind (have

“objective reality” in it and for it, for the mind to be “as if of” things) is treated as a basic property, an unexplained explainer. But an adequate treatment of the representational dimension of discursive sapience should include an account both of representational *purport*, and of its *uptake*.

The topic to be investigated here, then, is intentionality in the sense of the propositional contentfulness of attitudes, not in the sense (if that should turn out to be different) of the directedness of sense. The aim is to understand ourselves as judges and agents, as concept-users who can reason both theoretically and practically. This is not to say that we should understand ourselves exclusively as sapient rather than sentient, in terms of understanding rather than awareness. ‘We’ is and by rights ought to be said in many ways. The point is just to register and delineate the way that is to be discussed here.

This inquiry is directed at the fanciest sort of intentionality, one that involves expressive capacities that cannot be made sense of apart from participation in linguistic practices. The aim is to offer sufficient conditions for a system of social practices to count as specifically *linguistic* practices, in the sense of defining an idiom that confers recognizably *propositional* contents on expressions, performances, and attitudes suitably caught up in those practices. Looking at this sort of high-grade intentionality accordingly risks being beastly to the beasts—not only by emphasizing sapience over sentience, comprehension over consciousness, but also by unfairly ignoring the sorts of beliefs and desires that are appropriately attributed to non- or pre-linguistic animals.

So it is a further criterion of adequacy of this explanatory enterprise that it have something to say about the lower grades of intentionality: not only as to how the lines should be drawn (corresponding to different senses of ‘we’), but also as to how the advent of the favored sort of linguistic intentionality can be made less mysterious. How can linguistic abilities arise out of nonlinguistic ones? Or to ask a related question, What would sentient creatures have to be able to do in order to count as sapient as well? What is needed is to tell a story about practices that are sufficient to confer propositionally contentful intentional states on those who engage in them, without presupposing such states on the part of the practitioners. The hope is that doing so will offer guidance concerning what would be involved in diagnosing aliens as exhibiting such states, and programming computers or teaching merely sentient animals to exhibit them.

## II. FROM INTENTIONAL STATE TO NORMATIVE STATUS

### 1. *Kant: Demarcation by Norms*

The demarcational proposal being pursued picks us out as the ones capable of judgment and action. Not only do we respond differentially to environing stimuli, we respond by forming perceptual judgments. Not

only do we produce behavior, we perform actions. Various ways of talking about this fundamental distinction have been put on the table. It can be made out in terms of *truth*. In perception what we do is responsively *take-true* some propositional content that is intelligible to us. In action what we do is responsively *make-true* some propositional content that is intelligible to us.

The distinction can be made out in terms of *reasons*. The judgments that are our perceptual responses to what is going on around us differ from responses that are not propositionally contentful (and so are not in that sense intelligible) in that they can serve as reasons, as premises from which further conclusions can be drawn. Actions, which alter what is going on around us in response to propositionally contentful intentions, differ from performances that are merely behavior (and so not intelligible in terms of the propositionally contentful intentions that elicit them) in that reasons can be given for them; they can appear as the conclusions of practical inferences.

The distinction can also be made out in terms of the employment of *concepts*. To be a perceiver rather than just an irritable organism is to be disposed to respond reliably and differentially to the perceptible environment by the application of appropriate concepts. To be an agent rather than just a behaver is to be disposed to respond reliably and differentially to applications of appropriate concepts by altering the accessible environment. Intelligibility in the sense of propositional contentfulness, whether the latter is conceived in terms of truth conditions or capacity to serve as a reason, is a matter of conceptual articulation—in the case of perception and action, that the reliably elicited response and the reliably eliciting stimulus, respectively, essentially involve the use of concepts.

So sapience, discursive intentionality, is concept-mongering. What is distinctive of specifically *conceptual* activity? Contemporary thought about the use of concepts owes great debts to Kant. One of his cardinal innovations is his introduction of the idea that conceptually structured activity is distinguished by its *normative* character. His fundamental insight is that judgments and actions are to be understood to begin with in terms of the special way in which we are *responsible* for them.

Kant understands concepts as having the form of *rules*, which is to say that they specify how something *ought* (according to the rule) to be done.<sup>1</sup> The understanding, the conceptual faculty, is the faculty of grasping rules—of appreciating the distinction between correct and incorrect application they determine. What is distinctive about judgments and doings—acts that have contents that one can take or make true and for which the demand for reasons is in order—is the way they are governed by rules. They are conceptually contentful and so are subject to evaluation according to the rules that express those contents. Being in an intentional state or performing an intentional action accordingly has a normative significance. It counts as undertaking (acquiring) an obligation or commitment; the content of the commitment is determined by the rules that are the concepts in terms of which the act or



state is articulated. Thus Kant's version of the sort of demarcation criterion being considered picks us out as distinctively *normative*, or rule-governed, creatures.

## 2. From Cartesian Certainty to Kantian Necessity

This emphasis on the normative significance of attributions of intentionally contentful states marks a decisive difference between Kantian and Cartesian ways of conceiving cognition and action. For Kant the important line is not that separating the mental and the material as two matter-of-factually different kinds of stuff. It is rather that separating what is subject to certain kinds of normative assessment and what is not. For Descartes, having a mind (grasping intentional contents) is having representings: states that purport or seem to represent something. Some things in the world exhibit this sort of property; others do not. Where Descartes puts forward a *descriptive* conception of intentionality, Kant puts forward a normative, or *prescriptive*, one—what matters is being the subject not of *properties* of a certain kind but of *proprieties* of a certain kind. The key to the conceptual is to be found not by investigating a special sort of mental *substance* that must be manipulated in applying concepts but by investigating the special sort of *authority* one becomes subject to in applying concepts—the way in which conceptually articulated acts are liable to assessments of correctness and incorrectness according to the concepts they involve.

This approach contrasts sharply with Cartesian demarcations of cognition and action according to the presence of items of a certain matter-of-factual kind. The objection is not to the details of Descartes's understanding of the descriptive features required for intentionally contentful states and acts: his conception of mental events that are self-intimating cognitions or infallibly performable volitions, takings-true and makings-true that are minimal in that they cannot fail to be successful. It is, more radically, that what sets off the intentional is its liability to assessments of correctness, its being subject to norms (which are understood as codified in rules), rather than any missing feature that it could be described as having or lacking.

Descartes inaugurated a new philosophical era by conceiving of what he took to be the *ontological* distinction between the mental and the physical in *epistemological* terms: in terms of accessibility to cognition—in terms, ultimately, of certainty. Kant launched a new philosophical epoch by shifting the center of concern from *certainty* to *necessity*. Where Descartes's descriptive conception of intentionality, centering on certainty, picks out as essential our grip on the concepts employed in cognition and action, Kant's normative conception of intentionality, centering on necessity, treats their grip on us as the heart of the matter. The attempt to understand the source, nature, and significance of the norms implicit in our concepts—both those that govern the theoretical employment of concepts in inquiry and knowl-

edge and those that govern their practical employment in deliberation and action—stands at the very center of Kant's philosophical enterprise. The most urgent question for Kant is how to understand the rulishness of concepts, how to understand their authority, bindingness, or validity. It is this normative character that he calls *Notwendigkeit* (necessity).

The nature and significance of the sea change from Cartesian certainty to Kantian necessity will be misunderstood unless it is kept in mind that by 'necessary' Kant means 'in accord with a rule'. It is in this sense that he is entitled to talk about the *natural* necessity whose recognition is implicit in cognitive or theoretical activity, and the *moral* necessity whose recognition is implicit in practical activity, as species of one genus. The key concept of each is obligation by a rule. It is tempting, but misleading, to understand Kant's use of the notion of necessity anachronistically, in terms of contemporary discussions of alethic modality. It is misleading because Kant's concerns are at base normative, in the sense that the fundamental categories are those of deontic modality, of commitment and entitlement, rather than of alethic modality, of necessity and possibility as those terms are used today. Kant's commitment to the primacy of the practical consists in seeing both theoretical and practical consciousness, cognitive and conative activity, in these ultimately normative terms.

So for Kant, concepts are to be understood by the theorist in terms of the *rules* that make them explicit, rules that specify how the concepts are *properly* or *correctly* applied and otherwise employed. Kant's appreciation of this normative significance of concept use is one of the lenses through which he views his relationship to his rationalist and empiricist predecessors. With the wisdom of hindsight, Kant can see a normative strand of concern with *responsibility* as fundamental to the Enlightenment.<sup>2</sup> Thus the *Meditations* is to be read as motivated by the demand that the meditator take personal responsibility for every claim officially endorsed—be prepared to answer for it, demonstrate entitlement to that commitment by justifying it. This theme remained merely implicit in Descartes's theorizing about us (as opposed to his motivation and methodology), for his explicit theory remains naturalistic (though not, of course, physicalistic).

Leibniz insists, against the empiricists, that inferential transitions between representations ought not to be assimilated to matter-of-factual, habitually acquired causal dispositions. He understands them rather as applications of general principles that must accordingly be available prior to any knowledge of empirical matters of fact. Kant takes over from his reading of Leibniz the general idea of rules as what underwrite cognitive assessments of inferences and judgments. He understands such a priori principles, however, not as very general statements of fact (even metaphysical fact), but as rules of reasoning. They are conceived not as descriptive but as prescriptive—as (in Sellars's phrase) "fraught with ought."

This lesson dovetails neatly with the moral he draws from Hume's

thought. On Kant's reading, Hume's contribution is to see that ordinary empirical discourse involves commitments that reach beyond the sequences of representations, however regular, in which the concepts deployed in that discourse are taken to originate. Kant's Hume recognizes that cognitive experience crucially involves the application and assessment of the correctness of the application of *rules*. For Kant, Hume's inquiry after the nature of the authority for this inferential extension takes the form of a quest for the nature of the necessity, understood as normative bindingness, exhibited by the rules implicit in empirical concepts. It is under this conception that Kant can assimilate Hume's point about the distinction between saying what *happens* (describing a regularity) and saying what is causally *necessary* (prescribing a rule) to his point about the distinction between saying what *is* and saying what *ought* to be. One need not buy the metaphysics that Kant uses to ground and explain his norms, nor accept his answer to Hume, in order to appreciate the transformation of perspective made possible by his emphasis on the normativeness of the conceptual, and hence of cognition and action—the latter distinguished in the first instance as what we are responsible for.

### 3. Frege: *Justification versus Causation*

Kant's lesson is taken over as a central theme by Frege, whose campaign against psychologism relies on respecting and enforcing the distinction between the normative significance of applying concepts and the causal consequences of doing so. For Frege, it is possible to investigate in a naturalistic way acts of judging or thinking (even thinking conceived in a dualistic way), but such an investigation inevitably overlooks the normative dimension that is essential to understanding the propositional contents that are judged or thought. Sometimes this point is put in terms of *reasons*, invoking inferential relations among judgeable contents, as when he complains that psychologism "loses the distinction between the grounds that justify a conviction and the causes that actually produce it,"<sup>3</sup> or again when he argues that "the laws in accordance with which we actually draw inferences are not to be identified with the laws of correct [*richtigen*] inference; otherwise we could never draw a wrong inference."<sup>4</sup> Sometimes the point is put in terms of *truth*, as when he says, "It is not the holding something to be true that concerns us, but the laws of truth. We can also think of these as prescriptions for making judgments; we must comply with them if our judgments are not to fail of the truth."<sup>5</sup> Put either way, the point is that concern with the propositional contents that are thought or judged is inseparable from the possibility of assessments of *correctness*. Besides empirical regularities, there are also proprieties governing inferring and holding-true. Besides questions of which judgeable contents *are* held true, and under what circumstances, there is the question of which ones *ought* to be, and when. Besides the question of what consequences holding-true or making a judgment with

a certain content in fact leads to, there is the question of what those consequences ought or must rationally be. Psychology can study the matter-of-factual *properties* of contentful acts of judging and inferring, but not the semantically determined *proprieties* that govern them, the norms according to which assessments of truth and rationality are to be made.

Psychologism misunderstands the pragmatic significance of semantic contents. It cannot make intelligible the applicability of norms governing the acts that exhibit them. The force of those acts is a prescriptive rather than a descriptive affair; apart from their liability to assessments of judgments as true and inferences as correct, there is no such thing as judgment or inference. To try to analyze the conceptual contents of judgments in terms of habits or dispositions governing sequences of brain states or mentalistically conceived ideas is to settle on the wrong sort of modality, on causal necessitation rather than rational or cognitive right. Such natural processes "are no more true than false; they are simply processes, as an eddy in the water is a process. And if we are to speak of a right, it can only be the right of things to happen as they do happen. One phantasm contradicts another no more than one eddy in water contradicts another."<sup>6</sup> Contradiction, correct inference, correct judgment are all *normative* notions, not *natural* ones.

The laws of nature do not forbid the making of contradictory judgments. Such judgments are forbidden in a normative sense. It is *incorrect* to endorse incompatible contents: rationally incorrect, incorrect according to rules of reason, prescriptions governing what is *proper* in the way of inferring and judging. The 'must' of justification or good inference is not the 'must' of causal compulsion. But the possibility of expressing each in terms of rules or laws, so central to Kant's enterprise, misleads if these two different sorts of laws are not kept distinct, as they are not by psychologism and associationism. "What makes us so prone to embrace such erroneous views is that we define the task of logic as the investigation of the laws of thought, whilst understanding by this expression something on the same footing as the laws of nature . . . So if we call them laws of thought, or, better, laws of judgment, we must not forget we are concerned here with laws which, like the principles of morals or the laws of the state, prescribe how we are to act, and do not, like the laws of nature, define the actual course of events."<sup>7</sup>

Frege expresses his views about the normative character of judgeable contents, which he understands as having truth conditions, and so about the application of concepts, which he understands as functions whose values are truth-values, by talking about the nature of logic, which he understands as the study of the laws of truth.

Logic, like ethics, can also be called a normative science.<sup>8</sup>

The property 'good' has a significance for the latter analogous to that which the property 'true' has for the former. Although our actions and

endeavours are all causally conditioned and explicable in psychological terms, they do not all deserve to be called good.<sup>9</sup> Discussion of just how these remarks about the normative or prescriptive character of logic relate to a commitment to the normative or prescriptive significance of the exhibition of conceptual content by judgments must await more detailed consideration of Frege's *Begriffsschrift* theory of logical vocabulary as expressive of conceptual contents, in Chapter 2.

#### 4. Wittgenstein on the Normative Significance of Intentional Content

Frege emphasizes that concern with the contents of concepts and judgments is inseparable from concern with the possibility of the concepts being correctly or incorrectly applied, the judgments correctly or incorrectly made, whether this correctness is conceived in terms of truth or of the goodness of inference. Understanding this point requires distinguishing normative from causal modalities. Beyond enforcing this distinction, however, Frege has little to say about the nature of the norms that matter for the study of conceptual contents, and so for logic—though of course he has a great deal to say about the structure of such contents (some of which will be rehearsed in subsequent chapters). His concerns are at base semantic rather than pragmatic. In the twentieth century, the great proponent of the thesis that *intentionally* contentful states and acts have an essentially *normative* pragmatic significance is the later Wittgenstein.

The starting point of his investigations is the insight that our ordinary understanding of states and acts of meaning, understanding, intending, or believing something is an understanding of them as states and acts that *commit* or *oblige* us to act and think in various ways. To perform its traditional role, the meaning of a linguistic expression must determine how it would be *correct* to use it in various contexts. To understand or grasp such a meaning is to be able to distinguish correct from incorrect uses. The view is not restricted to meaning and understanding but extends as well to such intentionally contentful states as believing and intending. This is one way of developing and extending Kant's point that to take what we do as judging and acting is to treat it as subject to certain kinds of assessments as to its correctness: truth (corresponding to the world) and success (corresponding to the intention). A particular belief may actually relate in various ways to how things are, but its content determines how it is *appropriate* for it to be related, according to the belief—namely that the content of the taking-true should be true. A particular intention may or may not settle how one *will* act, but its content determines how it is *appropriate* to act, according to the intention—namely by making-true that content. To say this is in no way to deny that occurrences of intentional states of meaning, understanding, intending, and believing have causal significances. It is simply to point out that

understanding them as contentful involves understanding them as also having normative significances.<sup>10</sup>

The issue constantly before us in Wittgenstein's later works is how to understand these normative significances of intentional contents—the way in which they incorporate standards for assessments of correctness. Many of his most characteristic lines of thought are explorations of the inaptness of thinking of the normative 'force', which determines how it would be appropriate to act, on the model of a special kind of causal 'force'. The sense in which understanding or grasping a meaning is the source of the *correct* use is quite different from the sense in which it is the source of what one in fact goes on to do.<sup>11</sup> Enforcing the Kantian and Fregean distinction between grounds in the order of justification and grounds in the order of causation is what is behind talk of the "hardness of the logical 'must'"<sup>12</sup> and the picture of the dominion or compulsion intentional states exercise over what counts as correct performance as a machine whose "super-rigid" construction precludes any sort of malfunction. "The machine as symbolizing its action . . . We talk as if these parts could only move in this way, as if they could not do anything else. How is this—do we forget the possibility of their bending, breaking off, melting, and so on? Yes; in many cases we don't think of that at all . . . And it is quite true: the movement of the machine-as-symbol is predetermined in a different sense from that in which the movement of any given actual machine is predetermined."<sup>13</sup>

The relation between the content of an intention and the performances that would fulfill that intention does not leave any room for misfire, corresponding to the melting or bending of the parts of a mechanism, for it is already a *normative* relation. The state is to settle what *ought* to be done, what *must* be done if it is to be realized. What actually does or would happen is another matter. The images of superrigidity—of being guided by rails that one cannot fall away from—are what one gets if one assimilates normative compulsion to causal compulsion, ignoring the Kantian distinction. That is, if the normative 'must' were a kind of causal 'must', it would have to be a puzzling, superrigid sort—but the point is not to start with this sort of naturalistic prejudice.

In fact, by contrast, "The laws of inference can be said to compel us; in the same sense, that is to say, as other laws in human society."<sup>14</sup> They determine, in a sense yet to be specified, what one *ought* to do. Being compelled in this sense is entirely compatible with failing to act as one 'must'. Indeed, the physical or causal possibility of making a mistake, or doing what one is obliged, by what one means, intends, believes, and desires, not to do, is essential to the conception of such states and shows the essentially normative nature of their significance. "'But I don't mean that what I do now [in grasping a sense] determines the future use causally and as a matter of experience, but that in a queer way, the use itself is already present'.—But of course it is, 'in some sense'! Really the only thing wrong with what you say

is the expression 'in a queer way'. The rest is all right; and the sentence only seems queer when one imagines a different language game for it from the one in which we actually use it"<sup>15</sup>—the different game, namely, of attributing *natural states* and properties, rather than *normative statuses* such as commitments. What is determined is not how one *will* act but how one *ought* to, given the sense or content grasped, or the rule one has endorsed. "How am I able to obey a rule?"—if this is not a question about causes, then it is about the justification for my following the rule in the way I do."<sup>16</sup> That is, it is a question about what actions accord with the rule, are obliged or permitted by it, rather than with what my grasp of it actually makes me do.

### 5. Norms and Intentional Explanation

Although Wittgenstein often uses specifically linguistic examples, and some commentators have focused exclusively on these cases, the normative phenomena he highlights are part and parcel of intentional attribution generally, whether or not language is in the picture. *Ceteris paribus*, one who believes that it is raining, and that moving under the tree is the only way to stay dry, and who desires to stay dry, *ought* to move under the tree. The intentional states make the action appropriate. Indeed, the concept of rationality achieves its paradigmatic application in just such circumstances, as conduct warranted by the attributed intentional states is characterized as *rationally* appropriate. The qualification marks the option being reserved to deny that the conduct is, say, morally, politically, or aesthetically appropriate. (It is a further question whether this explanatory role of rationality justifies conceiving of what is rationally appropriate as reducible to what is prudentially or instrumentally appropriate.) Taking the category of rationality to be essentially involved in intentional explanation, as Dennett and Davidson for instance do, is one way of recognizing the normative dimension of intentionality.

It is important to keep this acknowledgment distinct from further theses one may then want to endorse concerning that normative dimension. If one keeps one's eye resolutely on the causal dimension of intentional explanation, the normative aspect can be masked. For instance, Dennett conjoins his recognition of the constitutive role of rationality in intentional explanation with the claim that such explanation involves a substantive "rationality assumption," the assumption in effect that the system will by and large in fact act as it ought rationally to act. There is nothing wrong with considering explanations of this sort as intentional explanations, but it is important to distinguish normative intentional explanation from causal intentional explanation. The former explains only what the subject of the intentional states ought or is obliged or committed (rationally) to do in virtue of its exhibition of the attributed states. The latter makes the substantive rationality assumption and goes on to explain what in fact happens. Normative intentional

explanations are more fundamental; they are presupposed and built upon by causal ones.

The same normative considerations arise if one approaches intentional states from the direction of their functional roles in mediating perception and action. Where Wittgenstein, Sellars, and Dummett, for instance, look at meaning, understanding, believing, or intending something in terms of mastery of the public proprieties governing the use of linguistic expressions, others would see intentional states as already definable by their role in accounting for the conduct of rational agents, whether linguistically adept or not. Both views are functionalist, in a broad sense. They differ over how to draw the boundaries around the functional systems within which alone something can have the significance of an intentional state. The issue of the extent to which mastery of linguistic social practice is a prerequisite for possession of intentional states of various sorts is of course an important one. But one need not have settled on one or the other of these explanatory approaches in order to appreciate that intentional states belong on the normative side of the Kantian divide. For talk of functional roles is itself already normative talk.

Specifying the functional role of some state in a system is specifying how it *ought* to behave and interact with other states. It is with reference to such a role that one makes sense of the notion of a *malfunctioning* component, something that is not behaving as it is supposed to. One may go on to offer various stories about the source of the correctnesses involved in functional roles: invoking the intentions of the designer, the purposes of the user, or the way the system must function if it is to realize the evolutionary good of survival, or even the cognitive good of accurately representing its environment. The point is that all of these are accounts of the source of the norms of proper functioning that are an integral part of functional explanations. The job of a designer's drawing of a machine is to specify how the machine is supposed to work, how it ought to work, according to the intentions of the designer. It is for this reason that "we forget the possibility of [the pieces] bending, breaking off, melting, and so on." Wittgenstein is of course concerned to understand how it is possible to understand such normative roles. But the current point is just that the roles one seeks to specify, in explaining the significance of intentional states, must, to begin with, be understood in normative terms of proper or correct functioning. Once again, this is not to deny that the fact that some component or system ought (functionally) to behave in a certain way may under many circumstances have a causal significance regarding how it will in fact behave. The issues are in principle distinct, however, and causal functional accounts presuppose normative functional ones.

The recognition that the consequences of attributing intentionally contentful states must be specified in normative terms may be summed up in the slogan, "Attributing an intentional state is attributing a normative



status." This is one of the leading ideas to be pursued in the present investigation. Intentional states and acts have contents in virtue of which they are essentially liable to evaluations of the "force of the better reason." It is this mysterious "force"—evidently the core of the social practices of giving and asking for reasons—that Greek philosophy investigated and appealed to in demarcating us from the nonrational background of items that we can think and find out about but that cannot themselves think or find out about other things. This "force of the better reason" is a normative force. It concerns what further beliefs one is *committed* to acknowledge, what one *ought* to conclude, what one is *committed* or *entitled* to say or do. Talk of what is a reason for what has to do in the first instance not with how people do or *would* act but with how they *should* act, what they should acknowledge. The sophist may not in fact respond to this "force," but even the sophist *ought* to. To understand rationality and states whose contents are articulated according to their role in reasoning, one must understand the force of such 'ought's. The relevance of reasons to the attributing and undertaking of intentional states and acts is *prima facie* reason to employ a normative metalanguage in analyzing such activity.

The normative dimension of intentional attributions is equally apparent if the propositional contents of the states and acts that are attributed, exhibited, or performed are conceived, not in terms of their accessibility to *reasons*, but in terms of there being circumstances under which they would be *true*. Assessments of truth, no less than assessments of rationality, are normative assessments. Truth and rationality are both forms of correctness. To ask whether a belief is true is to ask whether it is in some sense proper, just as to ask whether there are good reasons for it is to ask whether it is proper in a different sense. The business of truth talk is to evaluate the extent to which a state or act has fulfilled a certain kind of responsibility. This normative aspect of concern with truth can be masked by offering a descriptive, matter-of-factual account of what truth consists in. But doing so should be understood as offering a theory about this variety of semantic correctness, not as a denial that correctness is what is at issue. Thus Dummett argues that one does not understand the concept of truth when one has only a method for determining when it correctly applies to a claim or belief—a practical mastery of its circumstances of application. One must also know the *point* of applying it, must understand that truth is the *proper goal* of assertion and belief, that the language game of assertion and belief implicitly but essentially involves the injunction that one *ought* to speak and believe the truth. That is what one is supposed to be trying to do. Without an appreciation of this normative significance of application of the concept *truth*, one does not understand that concept.

Raising the question of what a belief or claim *represents* or is *about* can be understood as treating it as in a special way answerable for its *correctness* to what is represented, what it is about.<sup>17</sup> Thus the claim that semantically

or intentionally contentful states and acts have, as such, pragmatic significances that should be specified in normative terms does not depend upon what particular model (for instance, reasons, truth conditions, or representation) is employed in understanding such contents. The theoretical task of the intentional content of a state or act is to determine, in context, the normative significance of acquiring that state or performing that act: when it is appropriate or correct to do so and what the appropriate consequences of doing so are. The content is to determine proprieties of use, employment, or performance for states, acts, and expressions that exhibit or express such contents. The content must (in context) settle when it is *correct* to apply a concept in judging, believing, or claiming, and what correctly follows from such an application. Correctnesses of application are discussed under the general headings of assessments of truth or representation; correctnesses of inference are discussed under the general heading of assessments of rationality.<sup>18</sup> To pick out intentional states and acts as ones to which any of these sorts of assessments—truth, accuracy of representation, or reasonability—are in principle appropriate is to treat their normative articulation as essential to them. For this point, it does not matter which sort of assessment is treated as fundamental, whether the goodness of claiming of the sort concepts of truth try to capture, the goodness of representation that concepts of correspondence try to capture, or the goodness of reasoning of the sort concepts of rationality try to capture. All are *prima facie* normative or evaluative notions.

### III. FROM NORMS EXPLICIT IN RULES TO NORMS IMPLICIT IN PRACTICES

#### 1. *Regulism: Norms as Explicit Rules or Principles*

The first commitment being attributed to Wittgenstein, then, is to taking the significance of attributing intentional states to be normative, a matter of the difference it makes to the correctness or justification of possible performances (including the adoption of other intentional states). The second commitment he undertakes concerns how to understand the normative statuses of correct and incorrect, justified and not justified, which this approach to intentionality concentrates on. The question of how the normative significances of intentional states are to be taken to be related to the matter-of-factual consequences of those states, which would be one way into this issue, can be put to one side for the moment. It is a question Wittgenstein is much interested in, but it ought to be seen as arising at a different point in the argument. For an account of the normative pole of the Kantian dualism need not take the form of a specification of how the normative is related to the nonnormative. Instead, Wittgenstein considers, and rejects, a particular model of correctness and incorrectness, roughly Kant's, in which what makes

a performance correct or not is its relation to some explicit rule. To understand his argument and the lesson he draws from it, it is necessary to see what this model of the normative is, and for what sort of explanatory role he claims it is unsuitable.

According to this more specific Kantian view,<sup>19</sup> norms just *are* rules of conduct. Normative assessments of performances are understood as always having the form of assessments of the extent to which those performances accord with some rule. Reference to proprieties of performance is taken as indirect reference to rules, which determine what is proper by explicitly *saying* what is proper. On this account, acts are liable to normative assessments insofar as they are governed by propositionally *explicit* prescriptions, prohibitions, and permissions. These may be conceived as rules, or alternatively as principles, laws, commands, contracts, or conventions. Each of these determines what one may or must do by *saying* what one may or must do. For a performance to be correct is, on this model, for the rules to permit or require it, for it to be in accord with principle, for the law to allow or demand it, for it to be commanded or contracted. It is because Kant is someone for whom the normative always appears in the explicit form of rules, laws, and commandments that he could see the rationalists' insistence on the essential role of principles in cognition and action as a dark appreciation of the fundamentally normative character of those faculties. It is for this reason that when Kant wants to say that we are creatures distinguished from others by the normative dimension of our conduct (both cognitive and practical), he puts this in terms of our being bound by *rules*.

On an approach according to which normative assessment of conduct—whether prospectively, in deliberation, or retrospectively, in appraisal—always begins with the question of what rule is followed in producing the performances in question, norms are likened to laws in the sense of statutes. For conduct is legally appropriate or inappropriate just insofar as it is governed by some explicit law that says it is. Assessments of legal praise and blame must at least implicitly appeal to the relation of the performance in question to some law. In this way, the model appeals to a familiar institutional context, in which the norms most in evidence clearly take the form of explicit principles, commands, and the like.

The influence of the jurisprudential analogy is evident in Kant's conception of the normative aspect of cognition and action in terms of following rules. Kant inherits the Enlightenment tradition, handed down from Grotius and Pufendorf, which first studied the normative in the form of positive and natural laws, conceived as the explicit commandments of sovereigns or superiors of one sort or another. As a result, Kant takes it for granted that it is appropriate to call a 'rule' or a 'law' whatever it is that determines the propriety or impropriety of some judgment or performance. For him, as for most philosophers before this century, explicit rules and principles are not simply one form among others that the normative might assume. Rules are

the form of the norm as such. This view, that proprieties of *practice* are always and everywhere to be conceived as expressions of the bindingness of underlying *principles*, may be called *regulism* about norms.<sup>20</sup>

According to this intellectualist, platonist conception of norms, common to Kant and Frege, to assess correctness is always to make at least implicit reference to a rule or principle that *determines* what is correct by explicitly *saying* so. In the best-known portion of his discussion of rule-following in the *Philosophical Investigations*, Wittgenstein argues that proprieties of performance that are governed by explicit rules do not form an autonomous stratum of normative statuses, one that could exist though no other did. Rather, proprieties governed by explicit rules rest on proprieties governed by practice. Norms that are *explicit* in the form of rules presuppose norms *implicit* in practices.

## 2. Wittgenstein's Regress Argument

Norms explicit as rules presuppose norms implicit in practices because a rule specifying how something is correctly done (how a word ought to be used, how a piano ought to be tuned) must be applied to particular circumstances, and applying a rule in particular circumstances is itself essentially something that can be done correctly or incorrectly. A rule, principle, or command has normative significance for performances only in the context of practices determining how it is correctly applied. For any particular performance and any rule, there will be ways of applying the rule so as to forbid the performance, and ways of applying it so as to permit or require it. The rule determines proprieties of performance only when correctly applied.

If correctnesses of performance are determined by rules only against the background of correctnesses of application of the rule, how are these latter correctnesses to be understood? If the regulist understanding of all norms as rules is right, then applications of a rule should themselves be understood as correct insofar as they accord with some further rule. Only if this is so can the rule-conception play the explanatory role of being the model for understanding *all* norms. A rule for applying a rule Wittgenstein calls an "interpretation" (*Deutung*). "There is an inclination to say: every action according to the rule is an interpretation. But we ought to restrict the term 'interpretation' to the substitution of one expression of the rule for another."<sup>21</sup> The question of the autonomy of the intellectualist conception of norms, presupposed by the claim that rules are the form of the normative, is the question of whether the normative can be understood as "rules all the way down," or whether rule-like proprieties depend on some more primitive sort of *practical* propriety. Wittgenstein argues that the latter is the case. Rules do not apply themselves; they determine correctnesses of performance only in the context of practices of distinguishing correct from incorrect applications of the rules. To conceive these practical proprieties of application as themselves rule-governed is to embark on a regress. Sooner or later the theorist will have to

acknowledge the existence of practical distinctions between what is appropriate and what not, admitting appropriatenesses according to practice as well as according to rules or explicit principles.

This regress argument shows that the platonist conception of norms as rules is not an autonomous one, and so does not describe the fundamental form of norm. "What does a game look like that is everywhere bounded by rules? whose rules never let a doubt creep in, but stop up all the cracks where it might?—Can't we imagine a rule determining the application of a rule, and a doubt which it removes—and so on?"<sup>22</sup> In each case the doubt is the possibility of a mistake, of going wrong, of acting incorrectly, for instance in applying a rule. The point is to be that a rule can remove such a doubt, settle what is correct to do, only insofar as it is itself correctly applied. "But how can a rule show me what I have to do at this point? Whatever I do is, on some interpretation, in accord with the rule.'—That is not what we ought to say, but rather: any interpretation [*Deutung*] still hangs in the air along with what it interprets [*dem Gedeuteten*], and cannot give it any support. Interpretations by themselves do not determine meaning."<sup>23</sup> No sequence of interpretations can eliminate the need to apply the final rules, and this is always itself subject to normative assessment. Applied incorrectly, any interpretation misleads. The rule says how to do one thing correctly only on the assumption that one can do something else correctly, namely apply the rule.

This was our paradox: no course of action could be determined by a rule, because every course of action can be made out to accord with the rule. The answer was: if everything can be made out to accord with the rule, then it can also be made out to conflict with it. And so there would be neither accord nor conflict here. It can be seen that there is a misunderstanding here from the mere fact that in the course of our argument we give one interpretation after another; as if each one contented us for at least a moment, until we thought of yet another standing behind it. What this shows is that there is a way of grasping a rule [*eine Auffassung einer Regel*] which is not an interpretation, but which is exhibited in what we call "obeying the rule" and "going against it" in actual cases.<sup>24</sup>

Absent such a practical way of grasping norms, no sense can be made of the distinction between correct and incorrect performance—of the difference between acting according to the norm and acting against it. Norms would then be unintelligible.

### 3. Wittgenstein's Pragmatism about Norms

The conclusion of the regress argument is that there is a need for a *pragmatist* conception of norms—a notion of primitive correctnesses of performance *implicit* in *practice* that precede and are presupposed by their *explicit* formulation in *rules* and *principles*. "To use the word without a

justification does not mean to use it wrongfully [*zu Unrecht gebrauchen*]."<sup>25</sup> There is a kind of correctness that does not depend on explicit justifications, a kind of correctness of practice.

And hence also 'obeying a rule' is a practice [*Praxis*].<sup>26</sup>

—To obey a rule, to make a report, to give an order, to play a game of chess, are customs [*Gepflogenheiten*] (uses, institutions).<sup>27</sup>

The regress argument does not by itself provide such a conception of proprieties of practice; it just shows that without one we cannot understand how rules can codify the correctnesses that they do.

This argument shares its form with the regress Lewis Carroll invokes in "The Tortoise and Achilles"<sup>28</sup> but takes that line of thought one level deeper. That story depends on the fact that in a formal logical system, statements are inferentially inert. Even conditionals, whose expressive job it is to make inferential relations explicit as the contents of claims, license inferential transitions from premises to conclusions only in the context of rules permitting detachment. Rules are needed to give claims, even conditional claims, a normative significance for action. Rules specify how conditionals are to be used—how it would be correct to use them. It is the rules that fix the inference-licensing role of conditionals, and so their significance for what it is correct to do (infer, assert). Although particular rules can be traded in for axioms (in the form of conditional claims), one cannot in principle trade in *all* rules for axioms. So one cannot express all of the rules that govern inferences in a logical system in the form of propositionally explicit postulates within that system.

Carroll uses the regress of conditionals that results from the attempt to replace the rule of conditional detachment by explicitly postulated conditionals as an argument to show this. Wittgenstein's regress-of-rules argument shows further that, while rules can codify the pragmatic normative significance of claims, they do so only against a background of practices permitting the distinguishing of correct from incorrect applications of those rules. Carroll's point is that the significance of claims for what it is correct to do must somehow be secured. Logical claims, like others, must have some normative pragmatic significance. Wittgenstein's point is then that conceiving such significances in regulist terms, as the invocation of rules of inference does, is not the whole story. Rule-based proprieties of performance depend on practice-based ones. The regulist, platonist, intellectualist conception of norms must be supplemented by that of the pragmatist.

Two commitments have now been attributed to Wittgenstein. The first is a *normative* thesis about the *pragmatics* of intentionality. The second is a *pragmatic* thesis about the *normativeness* of intentionality. In the first case, pragmatics is distinguished from semantics, as the theory of the significance of contentful states and performances from the theory of their contents. In

the second case, pragmatic theories of norms are distinguished from platonist theories, in treating as fundamental norms *implicit* in *practices* rather than norms *explicit* in *principles*. The first point enforces attention to the significance of intentional states for what it is *correct* to do. The second point is that proprieties of practice must be conceivable antecedently to their being expressly formulated into propositionally explicit governing rules or principles. For performances can be rule-governed only insofar as they are governed as well by practices of applying rules.

It is useful to approach the sort of understanding that is involved in mastering a practice, for instance a practice of applying or assessing applications of a rule, by means of Ryle's distinction between knowing *how* and knowing *that*.<sup>29</sup> Knowing how to do something is a matter of practical ability. To know how is just to be reliably able. Thus one knows how to ride a bicycle, apply a concept, draw an inference, and so on just in case one can discriminate in one's practice, in the performances one produces and assesses, between correct and incorrect ways of doing these things.

The explicit knowing-that corresponding to such implicit knowing-how is a theoretical formulation or expression of that practical ability, in a rule or principle, that *says* what is correct and what not. The intellectualist picture underwrites every bit of know-how with a bit of knowledge-that, which may be only implicit in practical discriminations. "Compare knowing and saying: how many feet high Mont Blanc is—how the word 'game' is used—how a clarinet sounds. If you are surprised that one can know something and not be able to say it, you are perhaps thinking of a case like the first. Certainly not like the third."<sup>30</sup> What Wittgenstein shows is that the intellectualist model will not do as an account of the nature of the normative as such. For when applied to the norms governing the application of rules and principles, it generates a regress, which can be halted only by acknowledging the existence of some more primitive form of norm. The regress is Wittgenstein's master argument for the appropriateness of the pragmatist, rather than the regulist-intellectualist, order of explanation.<sup>31</sup>

#### 4. *Sellars against Regulism*

Another thinker who, like Wittgenstein, takes his starting point from Kant's and Frege's appreciation of the normative character of intentionality (for him, coeval with language use) is Wilfrid Sellars. He takes up this theme in one of his earliest papers, published in 1947. The opening section of that paper is entitled "Behavior, Norm, and Semantic Meta-Language" and makes this point:

The psychologistic blunder with respect to 'means' is related to another fundamental error, that, namely, of confusing between (1) language as a descriptive category for which symbols are empirical classes to which

certain events belong (and hence are symbol-events) by virtue of performing an empirical function, with (2) language as an epistemological category for which the relation of type to token is not that of empirical class to member. . .

For the moment it will help clarify the epistemological distinction between symbol-types and symbol-tokens, if we think of the former as norms or standards, and of the latter as events which satisfy them. We can therefore, for the moment at least, contrast the above two senses of 'language' as the descriptive and the normative respectively. Making use of this distinction, we argue that 'meaning' or, better, 'designation' is a term belonging to language about languages in the second sense. Its primary employment is therefore in connection with linguistic expressions as norms.<sup>32</sup>

Like Wittgenstein, Sellars sees that an adequate conception of these norms must move beyond the pervasive regulist tradition, which can understand them only in the form of rules.

Focusing on linguistic intentionality, Sellars in another paper examines the regulist conception as it applies to the linguistic norms in virtue of which it is possible to say anything at all. "It seems plausible to say that a language is a system of expressions, the use of which is subject to certain rules. It would seem, thus, that learning to use a language is learning to obey the rules for the use of its expressions. However, taken as it stands, this thesis is subject to an obvious and devastating objection."<sup>33</sup> The objection is that taking 'correct' to mean 'correct according to a rule' generates a familiar sort of regress:

The refutation runs as follows: Thesis. Learning to use a language (*L*) is learning to obey the rules of *L*. But, a rule which enjoins the doing of an action (*A*) is a sentence in a language which contains an expression for *A*. Hence, a rule which enjoins the using of a linguistic expression (*E*) is a sentence in a language which contains an expression for *E*—in other words, a sentence in a metalanguage. Consequently, learning to obey the rules for *L* presupposes the ability to use the metalanguage (*ML*) in which the rules for *L* are formulated. So that, learning to use a language (*L*) presupposes having learned to use a metalanguage (*ML*). And by the same token, having learned to use *ML* presupposes having learned to use a metmetalanguage (*MML*) and so on. But, this is impossible (a vicious regress). Therefore, the thesis is absurd and must be rejected.<sup>34</sup>

The metalanguage expresses rules for the proper application of concepts of the object language. But these rules, too, must be applied. So the



metametalinguage expresses rules for applying the rules of the metalanguage, and so on.

If any talk is to be possible, there must be some meta . . . metalevel at which one has an understanding of rules that does not consist in offering another interpretation of them (according to rules formulated in a metalanguage) but which consists in being able to distinguish correct applications of the rule in practice. The question is how to understand such practical normative know-how. Although he, like Wittgenstein, uses 'rule' more broadly than is here recommended, Sellars is clearly after such a notion of norms implicit in practice: "We saw that a rule, properly speaking, isn't a rule unless it lives in behavior, rule-regulated behavior, even rule-violating behavior. Linguistically we always operate within a framework of living rules. (The snake which sheds one skin lives within another.) In attempting to grasp rules as rules from without, we are trying to have our cake and eat it. To describe rules is to describe the skeletons of rules. A rule is lived, not described."<sup>35</sup>

This line of thought, common to Wittgenstein and Sellars, raises the key question of how to understand proprieties of practice, without appealing to rules, interpretations, justifications, or other explicit claims that something is appropriate. What does the practical capacity or 'know-how' to distinguish correct from incorrect performances (for instance—but this is only one example—applications of a rule) consist in? This is to ask what it is to take or treat a performance as correct-according-to-a-practice. It should also be asked, What is it for an act to be correct-according-to-a-practice? Both questions are important ones to ask: In what sense can norms (proprieties, correctnesses) be implicit in a practice? and What is it for someone to acknowledge those implicit norms as governing or being binding on a range of performers or performances?

The answers to these questions may be more intimately related to one another than at first appears. To foreshadow: On the broadly phenomenalist line about norms that will be defended here, norms are in an important sense in the eye of the beholder, so that one cannot address the question of what implicit norms are, independently of the question of what it is to acknowledge them in practice. The direction of explanation to be pursued here first offers an account of the practical attitude of *taking* something to be correct-according-to-a-practice, and then explains the status of *being* correct-according-to-a-practice by appeal to those attitudes. Filling in a story about normative attitudes as assessments of normative status, and explaining how such attitudes are related both to those statuses and to what is actually done, will count as specifying a sense of "norms implicit in practice" just insofar as the result satisfies the criteria of adequacy imposed on the notion of practice by the regress-of-rules argument.

Another central explanatory criterion of adequacy for such a conception of implicit practical normative knowing-how is that it be possible in terms

of it to understand explicit knowing-*that*. The effect is to reverse the regulist-intellectualist order of explanation. The regulist starts with a notion of norms explicit in principles and is obliged then to develop an account of what it would be for such things to be implicit in practices. The pragmatist starts rather with a notion of norms implicit in practice and is obliged then to develop an account of what it would be for such things to become propositionally explicit, as claims or rules. One of the primary tasks of this book is accordingly to offer an account of what it is to take some propriety that is implicit in a practice and make it explicit in the form of a claim, principle, or rule.

### 5. *Regularism: Norms as Regularities*

The regress-of-rules or regress-of-interpretations argument common to Wittgenstein and Sellars sets up criteria of adequacy for an account of contentful states that acknowledges their essentially normative significance, their characteristic relevance to assessments of the correctness of acts (including the adoption of further states). It must be possible to make sense of a notion of norms implicit in practice—which participants in the practice are bound by, and can acknowledge being bound by—without appeal to any explicit rules or capacities on the part of those participants to understand and apply such rules. Since the regress arises when the rule-following model of being bound by norms is applied to the agent, one strategy for avoiding it is to shift to a different model. Perhaps rules are relevant only as describing regularities, and not as being followed in achieving them.

Sellars (who does not endorse it) introduces such an approach this way: "Now, at first sight there is a simple and straightforward way of preserving the essential claim of the thesis while freeing it from the refutation. It consists in substituting the phrase 'learning to conform to the rules . . .' for 'learning to obey the rules . . .' where 'conforming to a rule enjoining the doing of *A* in circumstances *C*' is to be equated simply with 'doing *A* when the circumstances are *C*'—regardless of how one comes to do it . . . A person who has the habit of doing *A* in *C* would then be conforming to the above rule even though the idea that he was to do *A* in *C* never occurred to him, and even though he had no language for referring to either *A* or *C*."<sup>36</sup> What generates the regress is the demand that each practical capacity to act appropriately be analyzed as following an explicit rule that *says* what is appropriate, since understanding what is said by such a rule turns out to involve further practical mastery of proprieties.

If the practices in which norms are implicit are understood simply as regularities of performance, then there is nothing the practitioner need already understand. If such regularities of performance can be treated as practices governed by implicit norms, then there will be no regress or circularity in appealing to them as part of an account of knowing-*that*, of expressing

norms explicitly in rules and principles. For the only one who needs to understand how to apply correctly the rule conforming to which makes performances count as regular is the theorist who describes the regularity in terms of that rule. The norms implicit in regularities of conduct can be expressed explicitly in rules, but need not be so expressible by those in whose regular conduct they are implicit.

The view that to talk about implicit norms is just to talk about regularities—that practices should be understood just as regularities of behavior—may be called the simple regularity theory. It is clear how such a *regularist* account of the normative avoids the regress that threatens *regulist* accounts. The proposal is to identify being correct according to (norms implicit in) practice—in the sense required to avoid the regress of rules as interpretations that plagues fully platonist accounts—with conforming to (norms explicit in) a rule, where ‘conforming to a rule’ is just producing performances that are regular in that they count (for us) as correct according to it. The immediate difficulty with such a proposal is that it threatens to obliterate the contrast between treating a performance as subject to normative assessment of some sort and treating it as subject to physical laws.

For this reason simple regularity theories seem to abandon the idea that the significance of contentful states is to be conceived in normative terms. No one doubts that actions and linguistic performances are subject to laws of the latter sort and so conform to rules or are regular. The thesis of the normative significance of intentional states sought to distinguish intentional states from states whose significance is merely causal, and that distinction seems to be taken back by the simple regularity account. After all, as Kant tells us, in this sense “everything in nature, in the inanimate as well as the animate world, happens according to rules . . . All nature is actually nothing but a nexus of appearances according to rules; and there is nothing without rules.”<sup>37</sup> Everything acts regularly, according to the laws of physics. In what special sense do intentional states then involve specifically normative significances?

For a regularist account to weather this challenge, it must be able to fund a distinction between what is in fact done and what ought to be done. It must make room for the permanent possibility of mistakes, for what is done or taken to be correct nonetheless to turn out to be incorrect or inappropriate, according to some rule or practice. The importance of this possibility to the genuinely normative character of the force or significance associated with contentful states is a central and striking theme in Wittgenstein’s later works. What is correct or appropriate, what is obligatory or permitted, what one is committed or entitled to do—these are normative matters. Without the distinction between what *is* done and what *ought* to be done, this insight is lost.

The simple regularity approach is committed to identifying the distinction between *correct* and *incorrect* performance with that between *regular* and

*irregular* performance. A norm implicit in a practice is just a pattern exhibited by behavior. To violate that norm, to make a mistake or act incorrectly according to that norm, is to break the pattern, to act irregularly. The progress promised by such a regularity account of proprieties of practice lies in the possibility of specifying the pattern or regularity in purely descriptive terms and then allowing the relation between regular and irregular performance to stand in for the normative distinction between what is correct and what is not. Wittgenstein explicitly considers and rejects this approach. Where his master argument against regulism has the form of an appeal to the regress of interpretations, his master argument against regularism has the form of an appeal to the possibility of gerrymandering.

The problem is that any particular set of performances exhibits many regularities. These will agree on the performances that have been produced and differ in their treatment of some possible performances that have not (yet) been produced. A performance can be denominated 'irregular' only with respect to a specified regularity, not *tout court*. Any further performance will count as regular with respect to some of the patterns exhibited by the original set and as irregular with respect to others. For anything one might go on to do, there is some regularity with respect to which it counts as "going on in the same way," continuing the previous pattern. Kripke has powerfully expounded the battery of arguments and examples that Wittgenstein brings to bear to establish the point in this connection.<sup>38</sup> There simply is no such thing as *the* pattern or regularity exhibited by a stretch of past behavior, which can be appealed to in judging some candidate bit of future behavior as regular or irregular, and hence, on this line, as correct or incorrect. For the simple regularist's identification of impropriety with irregularity to get a grip, it must be supplemented with some way of picking out, as somehow *privileged*, some out of all the regularities exhibited. To say this is to say that some regularities must be picked out as the ones that *ought* to be conformed to, some patterns as the ones that *ought* to be continued. The simple regularity view offers no suggestions as to how this might be done and therefore does not solve, but merely puts off, the question of how to understand the normative distinction between what is done and what ought to be done.

One might respond to the demand that there be some way to pick out the correct regularity, from all the descriptively adequate but incompatible candidates, by shifting what one describes, from finite sets of performances to the set of performances (for instance, applications of a concept) the individual is *disposed* to produce. This set is infinite, in that any bearer of an intentional state is disposed to respond, say by applying or refusing to apply the concept *red* or *prime*, in an infinite number of slightly different circumstances. Kripke argues that this appeal to dispositions nevertheless does not suffice to rule out regularities that agree in all the cases one has dispositions with respect to, and differ in others so remote (perhaps, in the case of *prime*, because the numbers involved are so large, and in the case of *red* because surrounding

circumstances are so peculiar) that one does not have dispositions to treat them one way rather than another.

This last argument is controversial, but it is not a controversy that need be entered into here; however it may be with the finiteness objection to a dispositional account of the regularities that, according to the line of thought being considered, are to play the role of norms implicit in practice, there is another more serious objection to it. No one ever acts incorrectly in the sense of violating his or her own dispositions. Indeed, to talk of 'violating' dispositions is illicitly to import normative vocabulary into a purely descriptive context. Understanding the norms implicit in practice as descriptively adequate rules codifying regularities of disposition (even if a unique set of such rules is forthcoming) loses the contrast between correct and mistaken performance that is of the essence of the sort of normative assessment being reconstructed. If whatever one is disposed to do counts for that reason as right, then the distinction of right and wrong, and so all normative force, has been lost. Thus the simple regularity view cannot be rescued from the gerrymandering objection by appealing to dispositions in order to single out or privilege a unique regularity.

The problem that Wittgenstein sets up, then, is to make sense of a notion of norms implicit in practice that will not lose either the notion of implicitness, as regulism does, or the notion of norms, as simple regularism does. McDowell puts the point nicely: "Wittgenstein's problem is to steer a course between a Scylla and a Charybdis. Scylla is the idea that understanding is always interpretation. We can avoid Scylla by stressing that, say, calling something 'green' can be like crying 'Help' when one is drowning—simply how one has learned to react to this situation. But then we risk steering on to Charybdis—the picture of a level at which there are no norms . . . How can a performance be nothing but a 'blind' reaction to a situation, not an attempt to act on interpretation (thus avoiding Scylla); and be a case of going by a rule (avoiding Charybdis)? The answer is: by belonging to a custom (*PI* 198), practice (*PI* 202), or institution (*RFM* VI-31)."<sup>39</sup> The Scylla of regulism is shown to be unacceptable by the regress-of-rules argument. The Charybdis of regularism is shown to be unacceptable by the gerrymandering-of-regularities argument.

If anything is to be made of the Kantian insight that there is a fundamental normative dimension to the application of concepts (and hence to the significance of discursive or propositionally contentful intentional states and performances), an account is needed of what it is for norms to be implicit in practices. Such practices must be construed both as not having to involve explicit rules and as distinct from mere regularities. Wittgenstein, the principled theoretical quietist, does not attempt to provide a theory of practices, nor would he endorse the project of doing so. The last thing he thinks we need is more philosophical theories. Nonetheless, one of the projects pursued in the rest of this work is to come up with an account of norms implicit in

practices that will satisfy the criteria of adequacy Wittgenstein's arguments have established.

#### IV. FROM NORMATIVE STATUS TO NORMATIVE ATTITUDE

##### 1. *Kant: Acting According to Conceptions of Rules*

Two theses have so far been attributed to Kant. First, the sort of intentionality characteristic of us, exhibited on the theoretical side in judgment and on the practical side in action, has an essential *normative* dimension. Second, norms are to be understood as having the form of *explicit rules*, or principles. The first of these has been endorsed, as expressing a fundamental insight. The second has been rejected, on the basis of Wittgenstein's argument from the regress of rules as interpretations of rules. The conclusion drawn was that norms that are explicitly expressed in the form of rules, which determine what is correct according to them by *saying* or describing what is correct, must be understood as only one form that norms can take. That form is intelligible only against a background that includes norms that are *implicit* in what is *done*, rather than *explicit* in what is *said*.<sup>40</sup> At least the norms involved in properly understanding what is said by rules, or indeed in properly understanding any explicit saying or thinking, must be construed as norms of practice, on pain of a vicious regress.

In Kant's account of us as normative creatures, however, these two theses are inseparably bound up with a third. As has already been pointed out, Kant takes it that everything in nature happens according to rules. Being subject to rules is not special to us discursive, that is concept-applying, subjects of judgment and action.<sup>41</sup> What is distinctive about us as normative creatures is the *way* in which we are subject to norms (for Kant, in the form of rules). As natural beings, we act according to rules. As rational beings, we act according to our *conceptions* of rules.<sup>42</sup> It is not being bound by necessity, acting according to rules, that sets us apart; it is being bound not just by natural but by rational necessity. Kant's whole practical philosophy, and in particular the second Critique, is devoted to offering an account of this distinction between two ways in which one can be related to rules. Most of the details of his way of working out this idea are special to the systematic philosophical setting he develops and inhabits and need not be rehearsed here. Two fundamental features of his idea, however, must be taken seriously by any attempt to pursue his point about the normative character of concept-users.

The first of these has already been remarked on in connection with Frege. It concerns the distinction between the causal modalities and the more properly normative 'ought's whose applicability to us is being considered as a criterion of demarcation. This is the phenomenon distinguishing the force of causal 'must's from the force of logical or rational 'must's that Wittgen-

stein invokes in connection with his discussion of misunderstandings of the 'hardness' of the latter in relation to the former. It is an essential feature of the sort of government by norms that Kant is pointing to that it is compatible with the possibility of *mistakes*, of those subject to the norms going *wrong*, *failing* to do what they are obliged by those norms to do, or doing what they are *not* entitled to do. The 'ought' involved in saying that a stone subject to no other forces *ought* to accelerate toward the center of the earth at a rate of 32 feet per second per second shows itself to have the force of an attribution of natural or causal necessity by entailing that the stone *will* so act. The claim that it in this sense ought to behave a certain way is incompatible with the claim that it does not do so. In contrast, no such entailment or incompatibility is involved in claims about how we intentional agents ought to behave, for instance what else one of us is committed to believe or to do by having beliefs and desires with particular contents. Leaving room for the possibility of mistakes and failures in this way is one of the essential distinguishing features of the 'ought's that express government by norms in the sense that is being taken as characteristic of us, as opposed to it. The sense in which we are *compelled* by the norms that matter for intentionality, norms dictating what we are under various circumstances obliged to believe and to do, is quite different from natural compulsion.

The second feature of Kant's idea addresses precisely the nature of this normative compulsion that is nevertheless compatible with recalcitrance. For he does not just distinguish the sense in which we are bound by these norms from the sense in which we are bound by natural necessity in the purely formal terms invoked by this familiar point about the possibility of our going wrong. He characterizes it substantively as acting according to a *conception* or a *representation* of a rule, rather than just according to a rule. Shorn of the details of his story about the nature of representations and the way they can affect what we do, the point he is making is that we act according to our *grasp* or *understanding* of rules. The rules do not immediately compel us, as natural ones do. Their compulsion is rather mediated by our *attitude* toward those rules. What makes us act as we do is not the rule or norm itself but our *acknowledgment* of it. It is the possibility of this intervening attitude that is missing in the relation between merely natural objects and the rules that govern them. The slippage possible in our acting according to our conception of a rule is made intelligible by distinguishing the sense in which one is bound by a rule whose grip on us depends on our recognition or acknowledgment of it as binding from the sense in which one can be bound by a rule whose grip does not depend on its being acknowledged. This explanatory strategy might be compared to Descartes's invocation of intervening representations in explaining the possibility of error about external things—though Kant need not be understood as following Descartes's path from an implicit appeal to the regress that threatens such representationalist pictures of cognition to a diagnosis of the relation be-

tween the subject and those mediating, error-enabling internal representations as itself immediate and hence immune to the possibility of error.

The idea underlying the demarcational strategy Kant introduces when he defines us, denizens of the realm of freedom, as beings that are capable of acting according to a conception of a rule by contrast to the denizens of the realm of nature, is that natural beings, who merely act according to rules, that is, regularly, are capable of acknowledging norms *only* by obedience to them. We rational beings are also capable of grasping or understanding the norms, of making assessments of correctness and incorrectness according to them. Those assessments play a role in determining what we go on to do—the phenomenon Kant denominates “the rational will.” But for us, in contrast to merely natural creatures, the assessment of the propriety of a performance is one thing, and the performance itself is another. The possibility of not doing what we nevertheless count as bound or obliged to do arises out of this distinction. What is special about us is the sort of grasp or uptake of normative significance that we are capable of. To be one of us in this sense of ‘us’ is to be the subject of normative *attitudes*, to be capable of acknowledging proprieties and improprieties of conduct, to be able to treat a performance as correct or incorrect.

## 2. Practical Normative Attitudes

It is a challenge to retain this insight about the significance of our normative attitudes while accommodating Wittgenstein’s pragmatist point about norms (and so jettisoning the intellectualist insistence on the explicitness of norms that colors Kant’s treatment). In order to do so, it must be possible to distinguish the attitude of acknowledging *implicitly* or *in practice* the correctness of some class of performances from merely exhibiting regularities of performance by producing only those that fall within that class. Otherwise, inanimate objects will count as acknowledging the correctness of laws of physics, and the distinction Kant points out is lost. As before, the challenge is to reject intellectualist regulism about norms without falling into nonnormative regularism.

Consideration of this third thesis of Kant’s sharpens the point, however, by focusing it on the capacity to adopt a normative practical attitude—to act in such a way as to attribute a normative significance, without doing so by *saying* that that is what one is doing. The question now becomes, What must one be able to *do* in order to count as *taking* or *treating* a performance as correct or incorrect? What is it for such a normative attitude—attributing a normative significance or status to a performance—to be implicit in practice? The importance of this question is a direct consequence of Kant’s point, once his rendering has been deintellectualized by replacing grasp of principles with mastery of practices.



It is a consequence of this criterion of adequacy that the practical performances that are assessments cannot be just the same performances that are assessed. Addressing the simplest case first, treating a performance as correct cannot be identified with producing it. For according to such an identification, the only way in which a norm can be acknowledged in practice is by obeying it, acting regularly according to it. But then it is impossible to treat performances as *incorrect*: the norms one counts as acknowledging simply are whatever regularities describe what one does. Such an account would collapse Kant's distinction between the way in which we are governed by norms we acknowledge and the way in which we are governed by natural laws independently of our acknowledgment of them.

Kant's principle that we are the ones who act not only according to rules but according to a conception of them is the claim that we are not merely *subject* to norms but *sensitive* to them. This principle has been taken over here by saying that we are characterized not only by normative *statuses*, but by normative *attitudes*—which is to say not only that our performances are correct or incorrect according to various rules but also that we can in our practice *treat* them as correct or incorrect according to various rules. Using 'assessment' to mean an assignment of normative significance—in the most basic case taking as correct or incorrect—the point may be put by saying that Kant's principle focuses demarcational interest on the normative *attitudes* exhibited by the activity of *assessing*, rather than just on the normative *statuses* being *assessed*. In order to respect the lessons of Wittgenstein's pragmatism about the normative, assessing must be understood as something *done*; the normative attitude must be construed as somehow implicit in the practice of the assessor, rather than explicit as the endorsement of a proposition. Construed in these practical terms, a consequence of Kant's distinction is that mere conformity to a norm is not even a candidate as a construal of the normative attitude of assessing conformity, which expresses the sort of sensitivity to norms that characterizes us.

For brutes or bits of the inanimate world to qualify as engaging in practices that implicitly acknowledge the applicability of norms, they would have to exhibit behavior that counts as treating conduct (their own or that of others) *as* correct or incorrect. Of course such things do respond differentially to their own and each other's antics. The question is what role such a response must play in order to deserve to be called a practical taking or treating of some performance *as* correct or incorrect, perhaps in the way in which eating something deserves to be called a practical taking or treating of it as food. Any sort of reliable differential responsive disposition can be understood as inducing a classification of stimuli. Iron rusts in some environments and not others, and so can be interpreted as classifying its environments into two sorts, depending on which kind of response they tend to elicit. Such responsive classification is a primitive kind of practical taking *of* something *as*

something. It is in this sense that an animal's eating something can be interpreted as its thereby *taking* what it eats *as* food. The issue of current concern is what must be true of a behavioral response-kind for it to be *correct* or *appropriate* that something is taken as by being responded to in that way, rather than *food* or *wet*.

### 3. Sanctions

In part because much of the tradition of thought about normative status and attitudes has taken its departure from a legal model, it is natural to answer this question by invoking the notion of *sanctions*: of reward and punishment. According to such a *retributive* approach to assessment, one treats a performance as correct or appropriate by *rewarding* it, and as incorrect or inappropriate by *punishing* it. Such an account can take many forms, depending on how sanctions are construed. In the simplest case, applying a negative sanction might be understood in terms of corporal punishment; a prelinguistic community could express its practical grasp of a norm of conduct by beating with sticks any of its members who are perceived as transgressing that norm. In these terms it is possible to explain for instance what it is for there to be a practical norm in force according to which in order to be entitled to enter a particular hut one is obliged to display a leaf from a certain sort of tree. The communal response of beating anyone who attempts to enter without such a token gives leaves of the proper kind the normative significance, for the community members, of a license. In this way members of the community can show, by what they do, what they *take* to be appropriate and inappropriate conduct.

One example of this approach is Haugeland's account of practical norms in terms of social constellations of dispositions having a structure he calls "conformist." He asks us to imagine under this heading creatures who not only conform their behavior to that of other community members in the sense of imitating each other, and so tending to act alike (normally in the sense of typically) in similar circumstances, but also sanction each other's performances, making future behavior more likely to conform to ("cluster around") the emergent standards. "The clusters that coalesce can be called 'norms' (and not just groups or types) precisely because they are generated and maintained by censoriousness; the censure attendant on deviation automatically gives the standards (the extant clusters) a de facto normative force."<sup>43</sup> According to such an account the normative attitudes of taking or treating some performance as correct or incorrect are understood in terms of behavioral *reinforcement*, in the learning-theoretic sense. The advantage of such a way of putting things is that reinforcement is a purely functional descriptive notion, definable in abstraction from the particular considerations about familiar animals, in virtue of which beating them with a stick is

likely to function as negative reinforcement. Treating a performance as correct is taken to be positively sanctioning it, which is to say positively reinforcing it. Positively reinforcing a disposition to produce a performance of a certain kind as a response to a stimulus of a certain kind is responding to the response in such a way as to make it more likely in the future that a response of that kind will be elicited by a stimulus of the corresponding kind. Treating a performance as incorrect is taken to be negatively sanctioning it, which is to say negatively reinforcing it. Negatively reinforcing a disposition to produce a performance of a certain kind as a response to a stimulus of a certain kind is responding to the response in such a way as to make it less likely in the future that a response of that kind will be elicited by a stimulus of the corresponding kind.

The approach being considered distinguishes us as norm-governed creatures from merely regular natural creatures by the normative *attitudes* we evince—attitudes that express our grasp or practical conception of our behavior as governed by norms. These normative attitudes are understood in turn as *assessments*, assignments to performances of normative significance or status as correct or incorrect according to some norm. The assessing attitudes are then understood as dispositions to *sanction*, positively or negatively. Finally, sanctioning is understood in terms of *reinforcement*, which is a matter of the actual effect of the sanctioning or reinforcing responses on the responsive dispositions of the one whose performances are being reinforced, that is sanctioned, that is assessed.

Such a story is a kind of regularity theory, but not a simple regularity theory. It does not identify a norm wherever there is a regularity of behavior. In keeping with Kant's insight (as transposed from an intellectualist to a pragmatist key), norms are discerned only where attitudes—acknowledgments in practice of the bindingness of those norms—play a mediating role in regularities. Only insofar as regularities are brought about and sustained by effective assessments of propriety, in the form of responsive classifications of performances as correct or incorrect, are regularities taken to have specifically normative force. The possibility of incorrect, inappropriate, or mistaken performances—those that do not accord with the norm—is explicitly allowed for. Thus there is no danger of this sense of 'obligatory' collapsing into the sort of causal modality that governs merely natural happenings. Acknowledging a norm as in force is distinguishable in these terms from simply obeying it. A cardinal advantage of these theories is that while to this extent countenancing Kant's distinction between genuinely norm-governed and merely regular activity, they make intelligible how conduct that deserves to be called distinctively norm-governed could arise in the natural world.

The fundamental strategy pursued by such a theory is a promising one. As here elaborated, it involves three distinguishable commitments. First, Kant's distinction between acting according to a rule and acting according to a

conception of a rule is taken to express an important insight about the special way in which we are normative creatures. Second, the pragmatist regress-of-rules argument is taken to show that in order to make use of this insight, it is necessary that the sort of normative attitude that Kant takes to play an essential mediating role in our government by norms be understood as involving implicit acknowledgment of norms in practice. Specifically, it is necessary to make sense of the idea of practically taking or treating performances as correct or incorrect. Third, taking or treating performances as correct or incorrect, approving or disapproving them in practice, is explained in terms of positive and negative sanctions, rewards and punishments. This tripartite strategy is endorsed and pursued in the rest of this work. There are reasons not to be happy with the regularist way of working it out that has just been sketched, however.

Even the version of a regularity theory that Haugeland presents, which appeals only to patterns of positive and negative reinforcement to fund the notion of sanctions and thereby that of practical normative attitudes, merely puts off the issue of gerrymandering. Just as there is no such thing as *the* regularity of performance evinced by some actual course of conduct—because if there is one way of specifying it, there are an infinite number of distinguishable variants that overlap or agree about the specified performances and disagree about what counts as “going on in the same way”—so there is no such thing as *the* regularity that is being reinforced by a certain set of responses to responses, or even dispositions to respond to responses. The issue of gerrymandering, of how to privilege one specification of a regularity over equally qualified competitors, arises once more at the level of the reinforcing regularity. Again, simple regularity theories are subject to the objection that they conflate the categories of what is in fact done and what ought to be done, and hence that they fail to offer construals of genuinely normative significances of performances at all.

This is a way of failing to take sufficiently seriously Kant's distinction between acting according to a rule and according to a conception of a rule. Sanctions theories fund this crucial distinction by means of the distinction between producing a performance and assessing it. But assessing, sanctioning, is itself something that can be done correctly or incorrectly. If the normative status of being incorrect is to be understood in terms of the normative attitude of treating as incorrect by punishing, it seems that the identification required is not with the status of *actually* being punished but with that of *deserving* punishment, that is, being *correctly* punished. Of course sanctioners can be sanctioned in turn for their sanctioning, which is thereby treated as itself correctly or incorrectly done. Nonetheless, if actual reinforcement of dispositional regularities is all that is available to appeal to in making sense of this regress, it might still be claimed that what is instituted by this hierarchy of regularities of responses to regularities of responses ought not to count as genuinely normative.<sup>44</sup>

#### 4. Regularities of Communal Assessment

There are other dimensions along which it is instructive to consider sanctions theories of implicitly normative practice. Accounts that endorse the tripartite strategy by rendering practical acknowledgment of normative significance in terms of reward and punishment understand normative *status*, paradigmatically the significance of a performance as correct or incorrect, in terms of normative *attitudes*, paradigmatically *taking* or *treating as* correct or incorrect. Their characteristic suggestion is that this sort of practical attitude of assessment—responsive classification as correct or incorrect—is to be understood in terms of the practice of sanctioning. Giving pride of place in this way to normative attitudes in the understanding of normative statuses involves emphasizing the distinction of perspective between *assessing* a performance and *producing* that performance. Any theory that reconstructs Kant's distinction by appealing to this difference of perspectives on normative significance—the difference, namely, between consuming and producing normatively significant performances—is in one important sense a *social* theory of the sort of norm-governedness distinctive of us. Haugeland's censorious herd animals shape *each other's* behavior by their capacity not only to perform but to censure performance. Each animal in the community that is thereby constituted may (and perhaps to be a full-fledged member must) be able to do both, but as he conceives it, each act of censure involves two organisms, the censoring and the censured, the reinforcer and the reinforced.<sup>45</sup>

There is another sort of theory that combines the idea that normative statuses cannot be understood apart from normative attitudes with the idea that the relation between them involves social practices in a different but perhaps even more robust sense of 'social'. It has often been noticed that simple regularity theories of implicit norms gain no ground by shifting from concern with regularities characterizing the behavior of individuals to regularities characterizing the behavior of groups.<sup>46</sup> As a response to this concern, the leading idea of this sort of the construal of norms as implicit in social practices is that of *communal assessment*. On this approach, the key to the importance of the social is taken to lie in the possibility that the performances individual community members produce are assessed, responded to, or treated in practice as appropriate or inappropriate by the community to which the individual belongs. An individual might be taken implicitly to endorse or treat a performance as correct simply by producing it. The community, unlike the individual, need not be counted as having taken up a practical attitude regarding the propriety of the performance just in virtue of that performance's having been produced by one of its members. The class of performances produced by its members, rather, determines which fall within the scope of communal attitudes, which are liable to communal endorsement or repudiation.

Looking to assessments by the community does provide further resources for regularity theories. For a regularity can govern (that is, be displayed or conformed to by) the assessments of the community without each individual performance counting as correct according to it, and in that case a distinction between correct and incorrect performances by individuals is underwritten. In this way, the communal assessment theories are just like the theories already considered, with Haugeland's as an example, in which norms are taken to consist in regularities of practical appraisal. (It is irrelevant for the present point whether such appraisals are then understood in terms of sanctions, and the sanctions in terms of reinforcement.) The difference is that where the theories previously considered look to regularities of appraisals by individuals, the approach now on the table looks to regularities of appraisals by the community as a whole.<sup>47</sup>

There are two sorts of objections not yet considered to accounts that construe the norms implicit in practices in terms of regularities or dispositions regarding communal assessments or attributions of normative significance. First, the idea of communal performances, assessments, or verdicts on which it relies is a fiction. Second, the approach smuggles normative notions illicitly into what purports to be a reductive, nonnormative regularity theory. On the first point, communal assessment theorists have a tendency to personify the community, to talk about it as though it were able to do the same sorts of things that individual community members can do—perform additions, apply rules, assess performances, and so on. Thus to pick one page almost at random, Kripke talks about “the community's accepting” conditionals codifying relations between attributions of intentional states and commitments to act, what “the community regards as right,” what “the community endorses,” and so on.<sup>48</sup> This is a typical passage from Wright: “None of us can unilaterally make sense of the idea of correct employment of language save by reference to the authority of securable communal assent on the matter.”<sup>49</sup>

The difficulty with this way of talking is that assenting, endorsing, accepting, and regarding as right are in the first instance things done by *individuals*. It is not the community as such that assesses applications of the concept *yellow*, say, but individual members of that community. Any account that seeks to extend these notions to include cases where the subject or agent is a community should say explicitly what sort of performance or speech act is envisaged. Some communities have meetings, authorized representatives, or other ways of officially settling on a communal view or act, for instance of disapproval or endorsement. But this sort of thing is the exception and could in any case hardly be what is wanted for explaining either norms in general or conceptual, intentional, or linguistic norms more particularly.

This tendency to talk of the community as somehow having attitudes and producing performances of the sort more properly associated with individuals is neither accidental nor harmless. This *façon de parler* is of the essence of

the communal assessment approach. It is a manifestation of the orienting mistake (about which more will be heard later) of treating *I-we* relations rather than *I-thou* relations as the fundamental social structure.<sup>50</sup> Assessing, endorsing, and so on are all things we individuals do and attribute to each other, thereby constituting a community, a 'we'. But this insight is distorted by *I-we* spectacles—perhaps the same that have always been worn by political philosophers in conceiving their topic. The pretense of communal assessment is not harmless because the easy ways of cashing out the metaphor of community approval and so forth present familiar dilemmas. A notion of communal endorsement or repudiation might be built out of regularities of endorsement and repudiation by individual community members. But universal agreement is too much to ask, and how is it to be decided what less ought to count?

In fact, the approval of some almost always matters more to the community than that of others—though this division may be different for different issues. We recognize experts. But being an expert is having a certain *authority*, and that is a normative matter. One might go on to give an account of the status of having the authority to speak for the community on some matters in terms of the attitudes of attributing or recognizing that authority on the part of other community members. (A story along these lines is endorsed in what follows.) One might go on to offer sufficient conditions for the attributions of such attitudes, and so such statuses, to a community, entirely in nonnormative terms. (This move is repudiated in what follows.) But without some such story, how is one to understand talk of what the community endorses or repudiates?

Connected with this question is the problem of how community membership can be understood, in line with a communal assessment theory. This in turn is closely related to the previously mentioned problem, peculiar to regularity theories, of distinguishing those in whose practice the norms are implicit from those on whom those norms are binding. If 'norm implicit in a practice' is understood just as 'regularity truly descriptive of actual performances (or performances there are dispositions to produce under suitable circumstances)', and those performances are also the ones 'subject to' or 'governed by' the practices comprising them, then there is no possibility of irregularity, of violating a norm. Being a member of a community is rather being one who *ought* to conform to the norms implicit in the practices of the community. Community membership has this normative significance; it is a normative status.

Unless this status is understood in some way other than as being one who in fact exhibits the regularity in question, it will be impossible to violate a norm, because impossible to act irregularly. I cannot be out of step with any regularity that characterizes the behavior of each and every one of us. Extending that argument, if an account is to be offered of norms as social regularities, it is not by itself enough to identify what is correct for *all*

community members with what accords with the regularities descriptive of the practice of *some* of those community members (the experts). Some account is required of how those members are picked out. The distinction between the experts, the ones who have authority, whose actions and assessments cannot fail to be correct, and those who are subject to that authority, bound by the norms instituted by the regularities of performance of the experts, is a normative distinction, a distinction of normative status. Unless this distinction is itself understandable as a matter of regularities descriptive of the performance of community members, the social regularity account explains only some normative notions, while appealing to others.

Just as having the status of a community member is something with normative consequences—for one thereby is *subject* to a certain sort of authority, a certain standard of correctness—so having the status of an expert or official is something with normative consequences; for one thereby *exercises* a certain sort of authority. It is appropriate to ask whether the circumstances of application of the concept 'expert' (in the sense of authoritative assessor) could be specified in terms of regularities of conduct specified in nonnormative terms. Is there a distinction between actually assessing and being entitled to assess, one's assessments having authority? It is just in this sense that it is appropriate to ask more generally whether the circumstances of application of the concept 'community member' can be specified in terms of nonnormative regularities. Is there a distinction between actually being assessed and being properly assessed—being subject to the authority of assessments?

The point is that talk about the community to which a performer belongs is not obviously translatable into talk about regularities of individual performance. Belonging to the community is a concept used so as to have normative consequences of application, concerning the member's being responsible to the assessments of the community, being subject to its authority. An understanding of norms implicit in practice in terms of regularities of communal assessment requires the idea of a regularity that the performer is somehow bound by or answerable to. Exactly how must the performer be related to other performers to be appropriately taken to be responsible to their assessments?<sup>51</sup> It is possible that appeal will have to be made to something other than regularities of performance to secure this connection. Understanding normative status, including the normative status of being a community member, in terms of some sort of liability to being rewarded or punished is one thing; understanding that liability nonnormatively is another. The claim is not that these difficulties—with specifying what it is for a community to endorse or repudiate a performance, and the related questions of how to pick out the community members, those over whom such assessments have authority, and perhaps how to pick out experts whose assessments have such authority—are insuperable. These difficulties show only that the invocation of communal assessments does not by itself provide



a recipe for construing norms implicit in practices in terms of nonnormative regularities of performance (even assessing performances).

To recapitulate: Regularity theories of norms implicit in practice, that is of practical normative status, no matter whether they invoke specifically *social* regularities, or look in addition at regularities of *attitude* or assessment, or combine those moves in a theory that relies on regularities of communal assessment, are subject to two general sorts of objection. First, the gerrymandering argument challenges them to pick out a unique specification of the regularity in question, sufficient to project it so as to determine the application of the norms to novel cases. Second, bracketing those difficulties, attempts to define normative statuses in terms of nonnormative regularities can be criticized from two directions, as either failing to reconstruct some essential features of genuine normative statuses, or as covertly appealing to normative notions.<sup>52</sup> Under the first heading, they may fail to make room for the crucial distinction between *performance* and the *normative status* or significance of a performance, that is between what is done and what is correct or ought to be done, or they may fail to make room for the equally crucial distinction between the *normative status* of a performance and a *normative attitude* toward that status, that is between what is correct and what is taken to be correct. (More will be said about this second sort of insufficiency below.) The danger of regularity theories smuggling in normative notions arises both in specifying what it is for someone to be a member of a community in the sense of being governed by its norms, properly subject to assessments according to them, and again in specifying what it is for something to count as a communal attitude, an assessment by a community of conformity with its norms. This latter may take the form of specifying a subclass of community members whose assessments are imbued with communal authority and so have the status of official or expert assessments.

This mention of one sort of social practice theory is entirely preliminary. The issue of reasons for understanding the sort of practice in which norms are implicit as *social* practice, and in what sense of 'social', is addressed further along. It is one of the prime tasks of this work to elaborate a suitable notion of the social practices that institute the norms underlying explicitly propositional attitudes. For the approach eventually to be endorsed concerning the social nature of those implicit norms to be intelligible, something must be said here about the relation between normative and nonnormative vocabularies. Regularity theories, of whatever stripe, are (to adapt a phrase of Dretske's) attempts to bake a normative cake with nonnormative ingredients. Gerrymandering aside, the objections just mentioned represent the two ways such an enterprise can go wrong: by failing to produce a genuinely normative product or by employing some already normative raw materials. The discussion here of accounts that identify norms implicit in social practice with regularities of communal assessment arose from consideration of ways of pursuing the strategy that starts with Kant's distinction between

acting according to rules and acting according to conceptions of rules, and adapts it to accommodate Wittgenstein's pragmatic point about norms by rendering the normative attitude of taking or treating as correct (which according to the Kantian line is essential to the characteristic sort of normative status we exhibit) in terms of sanctions, of reward for what is (thereby) taken to be correct and punishment for what is (thereby) taken to be incorrect. The sanctions approach, the broadly retributive rendering of attitudes of assessment, offers aid and comfort to those concerned to tell normative stories with nonnormative vocabulary.

The animating idea is that the classification of performances needed to get this two-stage scheme off the ground, their being positive or negative sanctions, is something that can be made available in perfectly naturalistic terms. One might try to define the two sorts of sanctions in terms of the production of benefit versus harm to the one whose status or performance is being assessed. Or one might try to define them in terms of the preferences and aversions of the one assessed—respecting, as it were, the views of the sanctioned one rather than the sanctioner, as to what counts as benefiting and harming the sanctioned one. Less subjectively, reward and punishment might be understood in terms of giving pleasure and inflicting pain. Most austere, one might define rewarding or punishing a certain kind of behavior or performance in functional terms, as positively or negatively reinforcing the reliable dispositions to respond differentially to stimuli that are being realized by the performances responded to. This version takes it that to treat a response to a certain stimulus as incorrect is just to punish it, in the sense of responding to it in a way that in fact decreases the probability that the one being assessed will respond in that way to that sort of stimulus in the future.

### *5. Normative Sanctions*

In this connection it is important to realize that it is one thing to understand practical assessment as sanctioning, and quite another to understand sanctioning in nonnormative terms such as reinforcement. A retributive approach to the normative need not be given a naturalistic turn at all. Defining normative attitudes in terms of dispositions to apply sanctions does not by itself reduce the normative to the nonnormative—it just trades off one sort of norm for another. At the most basic level, to reward someone is to offer some good (either objectively or subjectively), and to punish them is conversely to inflict something bad. Benefit and harm, desirable and undesirable, are concepts that also have normative senses. Indeed, these senses would seem to be primary, so that some sort of reductive hypothesis would be needed to naturalize them. To turn the retributive story about normative attitudes into a naturalistic one, an account might for instance understand what is good (and so rewarding) in terms of what is desirable, what is desir-

able in terms of what is desired, and what is desired ultimately in terms of what is pursued.

Commitment to such a reduction is optional. Positive and negative sanctions may consist in acclaim and censure that itself has only a normative significance. A correct action might be rewarded by the grant of some extraordinary privilege or by release from some onerous obligation—and the status of such a response as reward need not depend on whether the one rewarded would in fact have been disposed to refrain from acting without the boon of entitlement or would in fact have been disposed to act so as to fulfill the obligation had it not been lifted. An incorrect action might be punished by withholding a license to act in certain other ways or by imposing an extraordinary obligation—and the status of such a response as punishment need not depend on whether the one punished is in fact disposed to refrain from acting even without the boon of entitlement, or is in fact disposed to act so as to fulfill the obligation imposed. In such cases one is rewarded or punished for what one does “in another world”—by a change in normative status rather than natural state.<sup>53</sup>

Consider once again the case of discerning a practical norm in force in a community, according to which to enter a particular hut one is obliged to display a leaf from a certain sort of tree. As pointed out above, the assessing response constituting the community's acknowledgment of such a norm (the attitude corresponding to the status) might in some cases be describable in nonnormative terms—one who violates the norm is beaten with sticks, the norm-violating behavior is negatively reinforced. But other cases are possible, for instance ones in which the assessing response is to punish by making other actions inappropriate—one who violates the norm is not permitted to attend the weekly festival. In such a case, the normative significance of transgression is itself specified in normative terms (of what is *appropriate*, of the transgressor is *entitled* to do). The punishment for violating one norm is an alteration in other normative statuses. Acting incorrectly alters what other performances are correct and incorrect.

Once again, it need not be assumed that the alteration of status according to which it becomes inappropriate to attend the tribal festival has the actual effect of disposing the transgressor not to attend. The alteration of status need itself have no reinforcing function. This could be so even if the assessing attitude corresponding to the consequential norm forbidding attendance at the festival is itself enforced by actual reinforcing responses—that is, even if it is the case that an attempt to attend the festival by one who is not entitled will be punished by beating the offending community member with sticks. In such a case, the norm regarding entitlement to attend the festival is intelligible in terms of attitudes expressed by sanctions specifiable in non-normative terms, while the norm regarding entitlement to enter the hut is intelligible in terms of attitudes expressed by sanctions specifiable only in

normative terms of consequences for entitlement to attend the festival. Since the norms governing festival attendance are directly nonnormatively intelligible, the norms governing hut entrance are also nonnormatively intelligible—but indirectly, at one remove. In this way one norm can depend on another, as the sanctions expressing assessments of the normative significance of performances according to the first norm consist in alterations of normative status with respect to the second norm.

If what qualifies some response to a performance as a sanction—and therefore, according to the retributive line being considered, as an assessment—is specifiable only in normative terms, that is in terms of the correctness or incorrectness, (the normative status) of further performances according to other norms, that kind of sanction can be thought of as being *internal* to the system of norms being discerned.<sup>54</sup> If, by contrast, what qualifies a response as a sanction is specifiable in wholly nonnormative terms of what various community members do or are disposed to do, without reference to the specifically normative status of their performances, that kind of sanction can be thought of as being *external* to the system of norms being discerned. In this terminology, the simple sort of dependence of one norm on another just considered occurs when the attitudes corresponding to one kind of normative status (e.g. propriety of hut entering) are expressed by normatively internal sanctions, defined in terms of another sort of normative status (e.g. propriety of festival attending). In the case described, that second sort of status is itself made intelligible by normatively external sanctions, ones that can be specified in terms of the movement of sticks and consequent alterations in dispositions to attempt festival attendance. But internal sanctions can be defined in terms of normative statuses that themselves are defined by internal sanctions referring one to still further norms. Clearly this sort of dependence of one norm on another according to the retributive paradigm can be extended and ramified, making sense of complex webs of interdependent normative statuses.

In the cases so far imagined, these webs of norms linked by internal sanctions are anchored, as each chain of definitional dependence terminates in some normative status that is definable independently, by external sanctions specified in nonnormative terms. Even this restriction can be relaxed. The consequences of an assessment of a performance as correct or incorrect with respect to one norm may extend no further than other assessments of correctness, with respect to other norms. It is possible to interpret a community as instituting normative statuses by their attitudes of assessment, even though each such status that is discerned is responded to by sanctions that involve only other normative statuses. It is compatible with the sanctions paradigm of assessment, and so of normative attitude, that it should be “norms all the way down.” Such an interpretation would not support any reduction of normative status to nonnormatively specifiable dispositions, whether to perform or to assess, whether individual or communal.

Norms acknowledged by external sanctions can be attributed to a community one by one, in an atomistic way. Attribution of any norm whose acknowledgment by the community takes the form of assessments expressed by internal sanctions, however, commits the interpreter to attributing also the norms on which it depends. Such dependences introduce a holistic element into the attribution of normative significances to the performances of a community. Using the retributive paradigm of normative attitudes of assessment to structure an interpretation of a community as exhibiting practices in which interdependent norms are implicit does not require that there be some nonnormatively specifiable behavior associated with the acknowledgment of each sort of normative status or significance discerned by that interpretation. An example of a system of practices in which the normative significances of performances must be attributed in a holistic way and are not translatable into nonnormatively specifiable dispositions is offered in Chapter 3, which presents sufficient conditions for such a system of practices to confer *propositional* contents on the statuses and attitudes it institutes.

A strategic divide looms here. Wittgenstein argued that an unproductive regress results from conceiving explicit rules as the only form of the normative. The lesson drawn from the regress of rules interpreting rules is the pragmatic one, that there must be "a way of grasping a rule which is not an interpretation, but which is exhibited in what we call 'obeying the rule' and 'going against it' in actual cases."<sup>55</sup> That is, there must be such a thing as norms that are implicit in practice. But what is the relation between such norm-laden practices and nonnormatively describable regularities of performance? The division of explanatory strategies arises over the question of whether the practices invoked to halt the regress<sup>56</sup> can be analyzed in terms of regularities and dispositions characterized without the use of normative vocabulary.

In line with Kant's insight that normative *attitudes*—the sort of uptake of or sensitivity to norms that he talked about in terms of *conceptions* of rules—are essential to the way in which our conduct is governed by norms, two suggestions have been put on the table.<sup>57</sup> The first is the idea of construing the normative attitude of taking or treating something as correct or incorrect in practice in terms of the application of positive and negative *sanctions*. The second is the idea that these assessing attitudes have a fundamentally social structure, so that the practices in which norms are implicit ought to be understood as essentially *social* practices. Each of these ideas could be given a reading in naturalistic or nonnormative terms, as part of a reductive explanatory strategy. Putting them together would then yield an approach that understands norms as implicit in regularities or dispositions regarding communal assessments of performances as correct or incorrect, and that understands such assessments in turn as behaviorally reinforcing rewards and punishments.

Wittgenstein certainly emphasizes the social nature of the practices un-

derlying the norms involved in discursive intentional states, and the importance to those practices of similarity or agreement of dispositions to respond to performances as correct or incorrect. There is accordingly a temptation to understand him as responding to the regress-of-rules argument by putting forward a reductive social regularity account of the practices in which norms are implicit.<sup>58</sup> But he need not be understood this way. His insistence that unless the responsive dispositions of a community are consilient, there can be no proprieties of practice is a point concerning presupposition, not reduction. Wittgenstein's is the somewhat delicate position first, that the usefulness of normative attributions, the viability of this stratum of discourse, presupposes a variety of regularities of performance and disposition; second, that those regularities obtain is *not* part of what is *asserted* by such attributions.

An analogy he comes back to again and again is the measurement of length, in which the possibility of practices of measurement presupposes features of the world such as the rigidity, spatial invariance under transportation, and temporal constancy of measuring rods, interpersonal comparability of measurements, the functional equivalence of various means of measuring the same length, the irrelevance to the result of such contextual features as whether the object measured is sacred or profane, to be used in sport or commerce, and so on. That we can be trained so as almost always to respond in the same way when applying concepts to novel cases, for instance, is a necessary condition of there being a practice determining what response is correct in such cases. But this is not to say that what it is for it to be correct consists in this agreement, as the reductive social regularity account of those norms would have it. There are three levels at which performances can be discussed: a level of norms explicit in rules and reasons, a level of norms implicit in practice, and a level of matter-of-factual regularities, individual and communal. To say that various claims made at the third level state necessary conditions for the applicability of vocabulary of the sort employed at the first two is not to make a reductive claim. The social regularity view conflates the second and the third levels, and so misunderstands Wittgenstein's remarks about the significance of matter-of-factual regularities, by taking them to involve commitment to the possibility of a reduction of the normative to the dispositional.<sup>59</sup>

## V. FROM ASSESSMENT TO THE SOCIAL INSTITUTION OF NORMS

### 1. *Pufendorf on the Institution of Norms by Attitudes*

As discursive beings whose characteristic activities are applying concepts, giving and asking for reasons, taking-true and making-true, we live and move and have our being in a space structured by norms. Yet we can describe, and largely successfully cope with, the not-us around us, while

restricting ourselves to a resolutely nonnormative vocabulary. In thinking about the relation between acting according to conceptions of rules (on a suitably pragmatic reading of what that consists in), as *we* do, and merely acting according to rules, as the rest of *it* does, it is important to distinguish two ways in which the normative significances we assign to things might be thought to be unnaturalized second-class citizens in an intrinsically insignificant natural world. These correspond to two different sorts of domestication to which normative statuses might be subjected. Couched in terms of supervenience, they are the claim that settling all the facts specifiable in *nonnormative* vocabulary settles all the facts specifiable in *normative* vocabulary, on the one hand, and the claim that settling all the facts concerning normative *attitudes* settles all the facts concerning normative *statuses*, on the other.

These are intimately related claims; the difference between them is subtle, and they are often run together. Each is the heir to a line of thought central to and characteristic of the Enlightenment project of disenchanting the natural world and humanizing values. The first can trace its origins to atoms-in-the-void physicalism—the conviction that a specification of the values of an appropriate range of dynamic variables for all the fundamental particles provides a *complete* description of everything that deserves to be called *real*. The second is animated by the humanistic thought that the merely natural world is devoid of values, that the worth of things and the fitness of actions is a product of our activity—that unlike natural properties, normative proprieties are in the eye of the human beholder. These ideas are of course at work in the thought of many Enlightenment philosophers. For present purposes it will suffice to consider briefly representative statements by one of the earliest.

The second line of thought emerges most clearly in the thought of the pioneering philosopher of law Samuel (Freiherr von) Pufendorf (1632–1694). Although not much read by philosophers today, his magnum opus *De Jure Naturae et Gentium* stands at the beginning of a tradition of Enlightenment thought about norms that culminates in Kant's practical philosophy (which was greatly influenced by Pufendorf). The relative unfamiliarity of these seminal views perhaps excuses quotation at greater than usual length. These passages all come from the opening of the work, in Chapter 1, entitled "On the Origin and Variety of Moral Entities": "It is for us to observe, how, chiefly for the direction of the acts of the will, a specific kind of attribute has been given to things and their natural motions, from which there has arisen a certain propriety in the actions of man, and that outstanding seemliness and order which adorn the life of men. Now these attributes are called Moral Entities, because by them the morals and actions of men are judged and tempered, so that they may attain a character and appearance different from the rude simplicity of dumb animals."<sup>60</sup> We are distinguished from the brutes by the fact that our actions are subject to assessment according to their

propriety, a special kind of attribute over and above the natural motions of things, an attribute that has somehow been *given* to them.

What is the source of those special normative attributes? "We seem able, accordingly, to define moral ideas most conveniently as certain modes [qualities], added to physical things or motions, by intelligent beings, primarily to direct and temper the freedom of the voluntary acts of man, and thereby to secure a certain orderliness and decorum in civilized life."<sup>61</sup> Intelligent beings add these properties to things by their activities. They are called 'moral' entities in virtue of their practical function as guides to action.

Now as the original way of producing physical entities is creation, so the way in which moral entities are produced can scarcely be better expressed than by the word *imposition*. For they do not arise out of the intrinsic nature of the physical properties of things, but they are super-added, at the will of intelligent entities, to things already existent and physically complete, and to their natural effects, and, indeed, come into existence only by the determination of their authors. And these authors give them also certain effects, which they can also remove at their own pleasure without any accompanying change in the object to which they had been added. Hence the active force which lies in them does not consist in their ability directly to produce any physical motion or change in any thing, but only in this, that it is made clear to men along what line they should govern their liberty of action.<sup>62</sup>

These norms are not part of the intrinsic nature of things, which is entirely indifferent to them. They are imposed by the will of intelligent beings and can affect things only through their effect on the will of such beings—beings who can act according to a conception of them. "Since, therefore, moral entities have been instituted to bring order into the lives of men, for which purpose it is required that they also, who must live according to their rule, should adopt a set standard in their relations toward one another, in determining their actions, and finally in fixing their attitude toward those things which are used in the lives of men; for this reason they are understood to be inherent primarily in men, but also in their actions, and even, to some extent, in things."<sup>63</sup> Our activity *institutes* norms, *imposes* normative significances on a natural world that is intrinsically without significance for the guidance or assessment of action. A normative significance is imposed on a nonnormative world, like a cloak thrown over its nakedness, by agents forming preferences, issuing orders, entering into agreements, praising and blaming, esteeming and assessing.<sup>64</sup>

One of the defining characteristics of early science is its disenchantment (*Entzauberung*, in the word we owe to Weber) of the world. The meanings and values that had previously been discerned in things are stripped off along with the supernatural and are understood as projections of human interests, concerns, and activities onto an essentially indifferent and insignificant mat-



ter. The Enlightenment disenchantment of the world and its assignment to us of responsibility for the norms, values, and significance we nonetheless find in the world are two sides of one coin. Meaningless objects and meaning-generating subjects are two aspects of one picture.<sup>65</sup> On this view, valuing is the source of values—a tradition carried on by figures as disparate as Mill and Nietzsche. Contractarian theories, and those that invoke positive law to explain various rights and obligations, are species of this genus. Each explains these deontic statuses in terms of what agents are doing in instituting or constitutively recognizing such entitlements and commitments.

Pufendorf does not suggest, and he does not believe, that the activity by which we institute norms is itself describable in the purely physical terms that suffice to describe the antics of merely natural objects. His claim is that the normative statuses of things, the normative significances we take them to have, are products of our practical normative attitudes, as expressed in our activity of imposing those significances and acknowledging them in assessments.<sup>66</sup> He does not conjoin this thesis with any sort of physicalism about the mechanism by which these moral secondary qualities arise from our practical activity. It is clearly possible to agree with the dictum of another Enlightenment thinker, Hamlet, that "There is nothing either good or bad, but thinking makes it so," without holding in addition a reductionist view about such thinking.

One who does conjoin these commitments is Hobbes. He explains, not good and evil, but *calling* things good and evil. He understands the use of the words 'good' and 'evil' as expressing appetites, desires, or aversions. He expresses his commitment to fundamental normative statuses being instituted by our attitudes this way: "But whatsoever is the object of any mans Appetite or Desire; that is it, which he for his part calleth *Good*: And the object of his Hate, and Aversion, *Evill* . . . For these words of Good, Evill . . . are ever used with relation to the person that useth them: There being nothing simply and absolutely so; nor any common Rule of Good and Evill, to be taken from the nature of the objects themselves; but from the Person of the man."<sup>67</sup> Gauthier comments on one difference between the sort of view Hobbes endorses concerning the relevant norm-instituting practical attitudes and the kind of view Hamlet endorses on this point: "If things considered in themselves are neither good nor bad, if there is no realm of value existing independently of animate beings and their activities, then thought is not the activity that summons value into being . . . Desire, not thought, and volition, not cognition, are the springs of good and evil."<sup>68</sup>

Pufendorf would not disagree. Where he does disagree is with Hobbes's subsequent endorsement of materialism about the will. The claim that normative statuses are instituted by our normative attitudes entails the claim that the normative proprieties so instituted are just natural properties of a special kind only in the context of a collateral claim that the norm-instituting practical attitudes can themselves be specified in nonnormative terms.

One can hold, as Pufendorf does, that there are no values apart from our acknowledgment or recognition of them, or more generally our attitudes toward them, without being thereby obliged in addition to understand those attitudes in terms of desires or preferences that can be characterized independently in value-free terms.

## 2. *Kantian Autonomy: The Authority of Norms Derives from Their Acknowledgment*

Pufendorf's idea that normative statuses are *instituted* by our practical attitudes makes a stronger claim than the idea previously extracted from Kant's demarcation of us as beings who act not only according to rules but according to our conceptions of rules. For the latter idea requires only that the normative statuses of demarcational interest essentially *involve* the uptake or grasp of such statuses, that is, our practical attitudes toward them. But normative statuses could be taken to be unintelligible apart from normative attitudes without thereby being taken to be instituted by and therefore in some sense to supervene on those attitudes. However Kant does in fact subscribe also to a version of the stronger thesis about attitudes instituting statuses, for the case of the genuinely moral normative statuses characteristic of us as agents.<sup>69</sup>

Kant's practical philosophy, his account of us as agents, takes its characteristic shape from his dual commitments to understanding us as *rational* and as *free*. To be rational, for him, means to be bound by rules. But Kant is concerned to reconcile our essential nature as in this way bound by norms with our radical autonomy. He combines the essential defining moment of our dependence on universals with that of our independence as particulars (as Hegel puts the point) in the thesis that the *authority* of these rules over us derives from our *acknowledgment* of them as binding on us. Our dignity as rational beings consists precisely in being bound only by rules we endorse, rules we have freely chosen (like Odysseus facing the Sirens) to bind ourselves with. We do not have the freedom to opt out entirely—choosing to be bound by no rules at all would be choosing to relinquish our rationality entirely. Yet if something other than our own attitudes and activity could bind us, we would not be free. Autonomy consists, as the etymology demands, in setting up laws for ourselves.

This view of Kant's inherits a venerable Enlightenment tradition.<sup>70</sup> It is based on a certain picture of the nature of the *authority* of rules or laws (the only form of norm considered). Pufendorf takes it that "since good repute, or moral necessity, and turpitude, are affections of human actions arising from their conformity to some norm or law, and law is the bidding of a superior, it does not appear that good repute or turpitude can be conceived to exist before law, and without the imposition of a superior."<sup>71</sup> More generally: "A

law may most conveniently be defined as a decree by which a superior obligates a subject to adapt his actions to the former's command."<sup>72</sup>

The key notion is that of a *superior*, someone who has *right* to command.

An obligation is properly laid on the mind of man by a superior.<sup>73</sup>

The power of obligating, that is, the faculty of laying an intrinsic necessity on persons to do something, properly lies in him who has authority or sovereignty.<sup>74</sup>

To see the question of authority in terms of who can command, make a rule binding, or lay down the law is another bit of fallout from the origin of thought about norms in thought about the institution of explicit positive laws. The issue of sovereignty is just the issue of "who's to be master, that's all," as the linguistic Leninist, Humpty Dumpty, says. The authority of norms depends on the nature of the author of the commands that make them explicit; their bindingness derives from the interpersonal relation of superior to subordinate.

The *consequences* of such a relation of authority being in force, what follows from a rule or law inheriting the authority of a superior lawgiver, Pufendorf conceives in terms of sanctions.<sup>75</sup> The *antecedents* or grounds on which such a relation is based, what makes one individual superior to a subordinate other in this normative sense, he conceives disjunctively: "Mere strength [to sanction] is not enough to lay an obligation on me at the desire of another, but that he should in addition have done me some special service, or I should of my own accord consent to his direction . . . But when a man of his own accord consents to the rule of another, he acknowledges by his own act that he must follow what he himself has decided."<sup>76</sup> The "special service" clause is put in to allow our creator a special claim on our obedience.

That special pleading aside, the standard Enlightenment thought, common as well to Hobbes and Rousseau among the progenitors of Kant's theory, is that our own *acknowledgment* or *endorsement* of a rule is the source of its authority over us—in short that our normative statuses such as obligation are instituted by our normative attitudes. Authority is not found in nature. The laws of nature do not bind us by obligation, but only by compulsion. The institution of authority is human work; we bind ourselves with norms. Contract theories are the result of combining a conflation of norms with their explicit expression in rules or laws, an understanding of their authority or bindingness on the hierarchical model of superior/subordinate (each given aid and comfort by the tradition of legalism), and an insistence on rational dignity as demanding autonomy.

Kant's reconciliation of us as free in virtue of being rational, with us as bound by norms in virtue of being rational—and so of freedom as constraint by a special kind of norm, the norms of rationality<sup>77</sup>—accordingly involves treating the normative status of moral obligation as instituted by normative

attitudes. It is our attitude toward a rule, our acknowledgment or recognition of moral necessity alone, that gives it a grip on us—not just in terms of its effect on our actual behavior, but in terms of our liability to assessment according to the rule that expresses that necessity. In this sense the norms that bind us rational creatures are instituted by our practical attitudes and activity. They are what we bring to the party. But while Kant in this way endorses the supervenience of moral normative status on moral normative attitude, he does not endorse any sort of naturalism or reductionism about those attitudes. He does not take it that specifications of those normative attitudes supervene on specifications of the movements of particles, described exclusively in the vocabulary of natural science. Grounding normative status in normative attitude does not entail relinquishing the distinction between normative proprieties and natural properties.

### 3. *Objectivity and the Social Institution of Conceptual Norms*

Kant also takes it that we are genuinely bound by the rules we endorse. This requires that once we endorse one, it is not up to us what it demands—there is some fact of the matter as to what we have thereby obliged ourselves to do. Although the status of being obliged to follow a particular rule is instituted by our attitudes, what is correct according to that rule is not simply determined by what we *take* to be correct according to it. The status of correctness of a performance according to a rule does not collapse into the attitude of assessing that performance as correct. Endorsing a rule gives it a grip on us. Part of that grip is that the rule does not mean just whatever we later might take it to mean. So Kant underwrites not only the possibility of *mistakes of performance*, which was already claimed to be essential to there being norms in play, but also the possibility of *mistakes of assessment*.

Wittgenstein appeals to this possibility as a criterion of adequacy for an account of norms being in force; talk of norms being implicit in a practice requires that there be room for a distinction between what is correct according to the norm and what the one whose performances are being assessed *takes* to be correct. He brings this consideration to bear against the possibility of setting up rules for oneself whose meaning is determined *only* by one's own actual dispositions to make assessments. This is the line of thought that concludes: "One would like to say: whatever is going to seem right to me is right. And that only means that here we can't talk about 'right'."<sup>78</sup> The thought here is that the distinction between status and assessment (the attitude of taking or treating a performance *as* correct, appropriate, or in order) is essential to the notion of genuinely normative status. It is motivated by the idea that assessing is itself something that can be done correctly or incorrectly, and furthermore that it is the norm according to which perform-

ances are being assessed that determines which assessments are correct or incorrect. If there is no distinction to be made between correct and incorrect assessments, then there is no sense in which the performances being assessed are governed by a norm *according to which* they are being assessed.

One of Wittgenstein's most important claims is that the practices in which the norms that articulate meanings and their uptake in understanding are implicit must be *social* practices. It is clear that the emphasis on their social character emerges for him somehow from the need to keep the notion of what one is committed to by the application of a concept distinct in principle from what one *takes* oneself to be committed to thereby. One natural way of understanding how such considerations can lead to the conclusion that discursive practice must be social practice is elaborated by Crispin Wright.<sup>79</sup> He combines an understanding of conceptual norms as instituted by practical normative attitudes—taking or treating various uses as correct or incorrect—with a way of maintaining a distinction between conceptual commitments and individual assessments of those commitments. He does so by identifying the normative status of being a correct application of a concept with being taken to be such a correct application, not by an individual, but by the whole community. According to this story, although individual performances can be correct or incorrect, and assessments of correctness by individuals can be correct or incorrect, no such difference applies to communal assessments. The community is incorrigible about what is a proper application of a concept and what is not. Communally endorsed applications of a rule, or, in the idiom employed here, acknowledgments of a norm implicit in the practice of the community, cannot be mistaken. "For the community itself there is no authority, so no standard to meet."<sup>80</sup>

Thus Wright secures a distinction between the commitments one undertakes in employing a particular concept and any individual's attitudes toward or assessments of those commitments, but at the cost of obliterating any such distinction between normative status and the attitudes of the whole community. There clearly are socially instituted norms of this sort. Whatever the Kwakiutl treat as an appropriate greeting gesture for their tribe, or a correctly constructed ceremonial hut, is one; it makes no sense to suppose that they could collectively be wrong about this sort of thing. The question is whether *conceptual* norms ought to be understood as being of this type. There is good reason to think they ought not. It is a fundamental feature of our understanding of our concepts that they incorporate *objective* commitments. Thus, our use of the term 'mass' is such that the facts settle whether the mass of the universe is large enough that it will eventually suffer gravitational collapse, independently of what we, even all of us and forever, take those facts to be. We could all be wrong in our assessment of this claim, could all be treating as a correct application of the concepts involved what is objectively an incorrect application of them.

On Wright's view the objectivity we take our conceptual norms to have is

an illusion that must be surrendered if they are to be properly understood. The normative attitudes discussed here under the heading of assessment—taking or treating applications of concepts as correct or incorrect—he talks about in terms of “ratification.” Understanding conceptual norms as objective, in the sense that the whole community can coherently be conceived to be wrong in its assessment of the commitments involved in some applications of its concepts, is taking them to be ratification-independent, in his terminology. McDowell, whose insistence on the objectivity of conceptual norms was discussed in this connection in Section IV, summarizes his disagreement with Wright on this point in this way: “In Wright’s reading . . . Wittgenstein’s point is that the natural contractual conception of understanding should not be discarded, but purged of the idea—which it must incorporate if the intuitive notion of objectivity is to have application—that the patterns to which our conceptions oblige us are ratification-independent. I expressed a suspicion above that this purging would not leave a residue recognizable as a conception of meaning and understanding at all.”<sup>81</sup> Wright takes it that understanding how the normative statuses involved in concept use are instituted by practical normative attitudes of assessing or ratifying the propriety of particular applications of concepts, while keeping normative statuses from collapsing into normative attitudes in a way that obliterates the norms entirely, at once requires understanding the practices of concept use and its assessment as *social* practices and relinquishing the idea that conceptual norms are *objective*.

A central aim of the present study is to show, by contrast, how these criteria of adequacy can be satisfied without giving up the objectivity of conceptual norms. Indeed the primary explanatory challenge to a social practice theory of discursive commitments is to show how, starting with the sort of norms for which Wright’s analysis is correct—normative statuses about which the community’s all-inclusive practical assessment cannot be mistaken, such as who is really married or what obligations are incurred by spitting in front of the chief—genuine, and therefore *objective*, conceptual norms can be elaborated. These bind the community of concept-users in such a way that it is possible not only for individuals but for the whole community to be mistaken in its assessments of what they require in particular cases.

How does objectivity precipitate out of the social soup of norms that are whatever the community takes them to be? According to the answer elaborated in Chapter 8, it is precisely the objectivity of conceptual norms, when properly understood, that leads to the requirement that the practices in which such norms are implicit be *social* practices. The objective representational dimension of conceptual content—the kind of correctness of claiming or concept application that answers not to individual or communal attitudes or assessments but to the properties of the things represented—turns out to depend on the *social* articulation of the inferential practice of giving and asking for reasons. Focusing on the distinction of social perspec-

tive between *acknowledging* (and thereby undertaking) a commitment oneself and *attributing* a commitment to another makes it possible to understand the objectivity of conceptual norms that consists in maintaining the distinction between the normative *statuses* they incorporate and the normative *attitudes* even of the whole community—while nonetheless understanding those statuses as instituted by the practical normative attitudes and assessments of community members. Far from precluding the possibility of conceptual objectivity, understanding the essentially social character of the discursive practice in which conceptual norms are implicit is just what makes such objectivity intelligible.

## VI. FROM INTENTIONAL INTERPRETATION TO ORIGINAL INTENTIONALITY

### 1. *The Stance Stance*

The normative house has many mansions. The particular norms of concern in this work are *discursive* normative statuses, the sort of commitment and entitlement that the use of concepts involves. These norms, it will be claimed, are instituted by *social* practices. These are practices that incorporate the distinction of social perspective between two kinds of practical attitude one can adopt toward a commitment: *acknowledging* it (oneself) and *attributing* it (to another). Elaborating an account along these lines is pursuing three of Wittgenstein's grand themes: the insistence on the *normative* character of language and intentionality, the *pragmatist* commitment to understanding these norms in terms of practices rather than exclusively in terms of rules, and the recognition of the essentially *social* character of such norms. One way in which the significance of the social character of the attitudes that institute intentional norms can begin to be approached is by considering the relation between the practical activity of intentional interpretation and the intentional states that are attributed by such interpretations.

Dennett's original account of intentional systems and intentional explanations provides a useful place to start.<sup>82</sup> One characteristic feature of that account is the idea that intentionality ought to be understood in terms of ascriptions of intentionality. Explanatory pride of place is granted to a certain sort of *attitude*, what Dennett calls a "stance." To adopt the intentional stance toward some system is to offer an intentional explanation of its behavior, by attributing intentional states to it. Adopting the intentional stance toward something is *taking* or treating it in practice *as* an intentional system. The status of being an intentional system, of exhibiting intentional states, is instituted by this attitude or stance: "A particular thing is an intentional system only in relation to the strategies of someone who is trying to explain and predict its behavior."<sup>83</sup> Dennett's explanatory strategy is first

to define what it is to adopt the intentional stance, that is to offer an intentional explanation, then to explain when it is appropriate to adopt that stance, and finally to define an intentional system as whatever is appropriately treated as one by adopting the intentional stance toward it. A few words are in order about each of these moves.

The beginning of wisdom about intentional explanation lies in appreciating the *normative* significance of attributing intentional states (mentioned in Section II of this chapter). Attributing suitably related beliefs and desires is attributing a certain sort of *reason* for action. Taking someone (1) to believe that it is raining and that the only way to stay dry is to open an umbrella and (2) to desire to stay dry is taking that individual to have a reason to open an umbrella.<sup>84</sup> To say this is not yet to say that the one who has such a reason *will* act according to it, even in the absence of competing reasons for incompatible courses of action. What follows immediately from the attribution of intentional states that amount to a reason for action is just that (*ceteris paribus*) the individual who has that reason *ought* to act in a certain way. This 'ought' is a *rational* ought—someone with those beliefs and those desires is rationally obliged or committed to act in a certain way. The significance of the states attributed is in the first instance a matter of the force of the better reason, rational force. That, as previously remarked, is a normative affair. Intentional interpretations attribute normative statuses, whose significance concerns practical proprieties. This is not to deny that reasons can be causes. It is just to unpack slightly what is meant by saying that they are reasons. The relation of such normative attributions of status and propriety to attributions of natural states and properties is a further issue.

Dennett acknowledges the normative core of intentional attribution and the corresponding distinction between physical and intentional explanation: "Deciding on the basis of available empirical evidence that something is a piece of copper or a lichen permits one to make predictions based on the empirical theories dealing with copper and lichens, but deciding on the basis of available evidence that something is (may be treated as) an intentional system permits predictions having a normative or a logical basis rather than an empirical one."<sup>85</sup> Attributing a natural state or property such as being copper supports *descriptive* conclusions about how the subject of those attributions *will* (in fact) behave. Attributing a normative status or propriety such as having beliefs and desires that amount to a reason for opening one's umbrella supports *prescriptive* conclusions about how the subject of those attributions *ought* (rationally) to behave. Within Dennett's project, however, the ultimate interest of intentional explanation lies in its use in deriving predictions concerning actual behavior. Some additional premise is required to get from the prescriptive conclusions that intentional attributions immediately supply to the descriptive predictions Dennett is concerned with.

He supplies the additional premise, in the form of a *substantive rationality assumption*, to the effect that agents generally do what one ought (ration-



ally) to do, what one is committed by one's intentional states to do. To be rational in Dennett's sense is to act as one rationally ought, to act as one's intentional states commit or oblige one to act. In order to derive predictions of actual behavior from attributions of intentional states, it is necessary to add the assumption that the subject to which those states are attributed is rational in this sense. In other words, intentional interpretation supplies a *primary* intentional explanation of the *normative* status of the one interpreted, an account of what performances are *appropriate* in the light of the beliefs and preferences attributed. Supplemented by a substantive rationality assumption, these normative characterizations can be used to ground predictions about actual performances, yielding a *secondary* intentional explanation of behavior described in nonnormative terms. The substantive rationality assumption provides the bridge that connects the normative significance of intentional attribution with the actual dispositions of the subject of such attribution.

Dennett's most controversial claim is his stance stance—his claim that there is no room for a distinction between actually being an intentional system and being appropriately treated as one. Intentional systems, things that have intentional states, just are whatever things it is predictively useful to adopt the intentional stance toward. The point of the stance idiom is that the notion of someone viewing or treating something as an intentional system is to be prior, in the order of explanation, to that of being an intentional system. The only notion of intentional system Dennett permits himself is "what one is treating something as when one offers intentional explanations of its behavior." Intentionally interpreting, adopting an intentional interpretive stance, is a practical attitude, and proprieties governing that practical attitude institute intentional states and hence normative statuses. Intentional states and intentional systems are, if not in the eye of the beholder, in the successful explanatory strategies of the theorist.

In the same way, Dennett distinguishes the significance of *ascribing* intentional states to a system from that of *describing* the system. Intentional ascriptions are appropriate according to their predictive utility, not their descriptive accuracy. The appeal to stances or attitudes as prior in the order of explanation to intentional states or normative statuses need not be read this way, however. For the contrast between talking about something in intentional vocabulary and talking about it in physical vocabulary is not for Dennett a distinction between adopting a stance and doing something else. The physical stance is also a stance. What appears in the instrumentalist reading as a distinction between what is really out there and what it is convenient from the point of view of prediction to attribute (the naturalized version of the proprieties of takings as intentional attributions), between describing and ascribing, between representing and adopting a stance or attitude, is, for Dennett, not itself a factual matter, a matter of what is out there, a matter of whether our representings do or do not correspond as they ought

to what is represented. It is rather a distinction between two stances that one may adopt. Taking there to be a physical fact of the matter determining the proprieties of our takings is adopting one stance; taking it that there are only predictive conveniences determining those proprieties is another. It is stances all the way down.

For Kant the difference between the realm of Nature and the realm of Freedom, and hence in the ordinary sense the distinction between facts and norms, is itself not a factual but a normative difference (the difference between acting according to rules and acting according to conceptions of rules). So one might say that for Dennett the difference between physical systems and intentional systems is itself a normative difference, a matter of the *propriety* of adopting different explanatory-predictive stances to the system in question. After all, for something to be a sample of copper is just for it to be proper or correct to treat it as one, in one sense of 'proper or correct' (the objective representational sense discussed at the end of the previous section), just as for something to be an intentional system is for it to be proper or correct to treat it as one, in another sense of 'proper or correct'. The question is how to understand the relation between the kinds of norms that govern the adoption of these different sorts of stance or attitude. It follows that Dennett's strategy of treating the normative significances of intentional states as instituted by the attitudes of interpreters does not by itself involve a commitment to reducing the normative to the nonnormative, insofar as it is *proprieties* of attitudes that are invoked. That reductive commitment comes in later, in explaining those proprieties. Understanding those proprieties in terms of predictive success, as Dennett does (a strategy different from that to be pursued here) gives an objective basis to the norms governing the adoption of the intentional stance. It puts Dennett in a position to say that talk of the predictive utility of adopting that stance is just a way—indeed, the only one available to us—of specifying an important kind of objective pattern of behavior. Thus the normative status of being an intentional system does not collapse into the adoption of normative attitudes of intentional interpretation.

## 2. *Different Stances and Kinds of Intentionality*

Understanding being an intentional system in terms of being appropriately taken or treated as an intentional system by being intentionally interpreted is not as such a reductive strategy for understanding intentionality in nonintentional terms. Offering intentional explanations of the behavior of others is something that only intentional systems can do. What is the relation between the intentionality that an intentional interpreter and attributor attributes, and that which the interpreter exhibits or possesses? To attribute beliefs (and desires and intentions), to adopt the intentional stance, one must have the concept of belief (desire, intention) and the capacity to

acquire beliefs concerning the appropriateness of applying that concept in understanding the behavior of candidate intentional systems. According to Dennett, intentional systems that can take up the intentional stance toward other systems have a special kind of intentionality. Intentional interpreters belong to "the subclass of intentional systems that have language, that can communicate": "Just as not all intentional systems currently known to us can fly or swim, so not all intentional systems can talk, but those which can do this raise special problems and opportunities when we come to ascribe beliefs and desires to them. That is a massive understatement, since without the talking intentional systems, of course, there would be no ascribing beliefs, no theorizing, no assuming rationality, no predicting."<sup>86</sup> Clearly, then, it is not possible to understand the second-class sort of intentionality attributed by creatures who offer intentional explanations of others, without understanding the first-class sort of intentionality those attributors themselves display. Dennett's assumption that possessing intentional concepts and attributing intentional states such as belief—that theorizing, predicting, assuming, and explaining—all presuppose specifically linguistic capacities is not universally shared, although good reasons for it will emerge. For present purposes, what matters is the distinction between first- and second-class intentionality—the kind possessed by attributors of intentionality, and the kind possessed by those to whom intentionality is attributed, rather than the specific characterization of the former. For it now appears that the intentionality of relatively simple systems such as animals and chess-playing computers, toward which Dennett takes it to be appropriate to adopt the intentional stance, can be understood only against the background of an understanding of the more complex systems capable of adopting that explanatory stance.

Thus it must be asked whether the fact that something is an intentional system in the first-class sense of attributing intentionality is a fact of the same general sort as the fact that something is an intentional system in the second-class sense of attributed intentionality. Dennett says of the second, attributed variety of intentionality, that the only facts in the vicinity are facts about the propriety of adopting a certain kind of stance toward it. Is the same thing true of the ascriber? Are the only facts about whether what one is doing is ascribing or attributing intentionality facts about the practical propriety of adopting a certain stance toward the interpreter, treating it in a certain way? Is adopting the intentional stance something one really does, or is the taking of a stance merely something that is sometimes appropriately attributed, so that it can be correct to adopt the stance that someone is adopting the intentional stance? Is it in this sense stances all the way down?

Distinguishing *simple intentional systems*, which are merely intentionally interpretable, from *interpreting intentional systems*, systems toward which the intentional stance can be adopted from systems that can adopt that stance toward others, is distinguishing instituting intentionality from instituted intentionality. Simple intentionality, which on this line is in the

eye of the beholder, is for that reason *dependent* on and in an important sense *derivative* from the intentionality exhibited by interpreters. The clearest examples of the derivative character of some intentionality or conceptual content are those in which interpreters explicitly assign some meaning to an intrinsically meaningless event by deciding to understand it in a certain way. They *make* an event mean something ("One if by land, and two if by sea . . .") by *taking* it to mean that, by understanding it that way. The meaning is conferred on the occurrence by the response to it that becomes appropriate, by the conclusions that are drawn from it (" . . . And I on the opposite shore will be"). The intentional content of the signal derives from the intentional content of the beliefs it makes appropriate for its audience. Noises and marks on paper do not mean anything all by themselves. Meaning is correlative with understanding, and they understand nothing. It is the possibility of *our* understanding them as expressing a content involving the application of concepts that makes *them* mean anything. Our understanding, our practices of interpretation institute that meaning, which derives from them.

The intentionality, the conceptual content, of noises and marks is borrowed from and dependent on that of the thoughts and beliefs that interpret them, the takings, or practical attitudes that attribute such content. On pain of an infinite regress, it seems necessary to distinguish the *derivative intentionality* such merely interpretable items display from the *original intentionality* their interpreters display. Clearly the simple intentionality of systems that can be interpreted as having and acting according to beliefs and desires is not derivative in the same sense in which that of inanimate marks and noises is. Nonetheless, on a view such as Dennett's the intentional contentfulness of the states of such systems depends on their interpretability by other, more capable systems. In the case both of inanimate and animate interpretables, the attempt to understand the sort of intentionality they display drives one back to the practical attitude or activity of interpreting. They exhibit no intentionality intelligible in its own right, apart from the practical attitudes of the interpreting systems.

Here, then, is a challenge: to maintain the stance toward both simple and interpreting intentional systems—that is, to acknowledge that the normative status of being such intentional systems is intelligible only by reference to the normative attitude of taking or treating something as such a system, that is interpreting it as one—while at the same time securing the distinction between original and derivative intentionality—and so not allowing the notion of intentional normative status to collapse into that of the normative attitude of intentional interpretation. This ought to seem hard to do. Indeed, Searle claims in effect that it is impossible<sup>87</sup>—that if derivative intentionality is to be intelligible, so must a sort of 'intrinsic' intentionality possessed by intentional interpreters, which can be made sense of quite apart from any reference to anyone taking or treating the states of those interpreters as intentionally contentful. From the point of view of the present project,

the regress argument he employs to derive that conclusion is flawed by its dependence on an inappropriate model of what it is to take or treat something as intentionally contentful. For following his treatment in *Speech Acts*,<sup>88</sup> he understands taking or treating a mark or noise as expressing a certain claim or proposition as depending on propositionally explicit beliefs and intentions regarding it—as interpreting it in Wittgenstein’s sense. A version of the regress-of-rules argument then shows that those beliefs and intentions cannot have their meaning conferred on them in the same way. But this leaves out the possibility of conferral of such content by *implicit* practical taking or treating of states, performances, and expressions as intentionally contentful. This is the possibility pursued in Chapter 3.

The theory developed in this work can be thought of as an account of the stance of attributing original intentionality. It offers an answer to the question, What features must one’s interpretation of a community exhibit in order properly to be said to be an interpretation of them as engaging in practices sufficient to confer genuinely propositional content on the performances, statuses, attitudes, and expressions caught up in those practices? The key to the account is that an interpretation of this sort must interpret community members as taking or treating each other in practice as adopting intentionally contentful commitments and other normative statuses. If the practices attributed to the community by the theorist have the right structure, then according to that interpretation, the community members’ practical attitudes institute normative statuses and confer intentional content on them; according to the interpretation, the intentional contentfulness of their states and performances is the product of their own activity, not that of the theorist interpreting that activity. Insofar as their intentionality is derivative—because the normative significance of their states is instituted by the attitudes adopted toward them—their intentionality derives from each other, not from outside the community. On this line, only communities, not individuals, can be interpreted as having original intentionality.

For this to work, the practices that institute the sort of normative status characteristic of intentional states must be *social* practices. Those practices essentially incorporate a distinction of social perspective between the attitudes of *undertaking* a commitment, as someone who believes that a bear is approaching might be taken to be committed to believing that an animal is approaching, and *attributing* a commitment, as the one who interprets another as having such a belief might do. The first sort of attitude toward a normative status must be attributed even to simple intentional systems—the rationality that is for Dennett the mother of intention is a way of talking about the sense of ‘ought’ in which one who believes a bear is approaching ought to believe that an animal is approaching. Just for that reason, the second sort of attitude is an implicit version of adopting the intentional stance. According to the account offered in Chapter 3 of the practices that confer distinctively propositional contents (and accordingly underlie all dis-

cursive intentionality, the conceptual contentfulness of expressions, performances, attitudes, and statuses), the practical normative attitudes of undertaking and attributing commitments come as a package—neither is intelligible apart from the other. Undertaking a commitment just is doing something that makes it appropriate for that commitment to be attributed. Normative statuses of the sort whose paradigm is provided by the inferentially articulated commitments constitutive of rationality are instituted by constellations of socially perspectival normative attitudes of attributing and undertaking such commitments. This is the *I-thou* structure of norm-instituting social practices that was contrasted above with the *I-we* sociality many theorists appeal to, and which is understood here as arising out of the more primitive perspectival variety.

### 3. Summary

The point of this chapter is to motivate the criteria of adequacy governing the account of discursive practice presented in Chapter 3, as well as the basic raw materials deployed there to satisfy those conditions. The first major point is the *normative* significance of intentional states, meanings, and the sort of understanding that is the uptake of those meanings. The second point is that norms that are *explicit* in the form of rules, principles, or claims (Wittgenstein's "interpretations") depend for their intelligibility—their determining a distinction between performances that are correct and incorrect, appropriate and inappropriate—on a more fundamental form of norms that are *implicit* in practice—in what is *done* rather than what is *said*. Making this distinction raises the question of how to understand the practice of making propositionally explicit claims (formulating principles, promulgating rules, and so on) in terms of norms that are implicit in practices.

The third point is that the attempt to understand norms implicit in practices by identifying the correct/incorrect distinction with the regular/irregular distinction (one strategy for reducing the normative to the nonnormative) will not work, for a reason parallel in form to the one that shows the need for a notion of norms implicit in practice in the first place. The regress-of-rules or regress-of-interpretations argument against *regularism* is that if rules were the only form of norms, they would fail to sort performances into those that are correct according to the rule and those that are not. For applying the rule is itself something that can be done correctly or incorrectly, and any performance that is correct according to one interpretation is incorrect according to others. The gerrymandering argument against *regularism* is that if norms are understood as regularities, they fail to sort performances into those that are correct (regular) and those that are not. Any course of conduct embodies many regularities, and any candidate performance that is regular according to one of them is irregular according to others. The two strategies

do not provide the resources to *privilege* one of the competing interpretations or regularities.

The fourth point, then, is that there is another move available for understanding what it is for norms to be implicit in practices. This is to look not just at what is *done*—the performances that might or might not accord with a norm (be appropriate or inappropriate)—but also at *assessments* of propriety. These are attitudes of taking or treating performances *as* correct or incorrect. If such attitudes are themselves understood on the model of propositionally explicit beliefs or commitments—as responding to a performance as correct by *saying* of it that it is correct—then the regress objection to regulism about norms reappears. But such assessing attitudes can also be understood as implicit in practice.

One way of doing that is to look to *sanctions*—treating a performance as correct by responding in practice with a reward (or the withholding of punishment) and treating it as incorrect by responding in practice with a punishment (or the withholding of a reward). What counts as a reward or punishment might be construed naturalistically, for instance as any response that positively or negatively reinforces the behavior responded to. Or it might be construed normatively, for instance in terms of the granting of special rights or the assignment of special obligations. Again, the assessing attitudes taken to be relevant to normative statuses can be taken to be implicit in the responses of other individuals, or of responses associated in some way with the whole community. In any of these cases, if the normative status of being a correct performance were identified solely by appeal to *regularities* exhibited by assessments, then the gerrymandering objection would be reinstated.

The fifth point, then, is that one way to demystify norms is to understand them as *instituted* by the practical attitudes of those who acknowledge them in their practice. Apart from such practical acknowledgment—taking or treating performances as correct or incorrect by responding to them as such in practice—performances have natural properties, but not normative proprieties; they cannot be understood as correct or incorrect without reference to their assessment or acknowledgment as such by those in whose practice the norms are implicit. At the same time, it must be acknowledged that a cardinal criterion of adequacy of any account of the *conceptual* norms implicit in discursive practice is that it make intelligible their *objectivity*. Doing so requires that the normative status of being a correct application of a concept not collapse into normative attitudes, as construing correctness as consisting just in being *taken* to be correct conflates them. The objectivity of conceptual norms requires that any attitude of taking, treating, or assessing as correct an application of a concept in forming a belief or making a claim be coherently conceivable as mistaken, because of how things are with the objects the belief or claim is about.

The next chapter opens the discussion of the propositional contents that are conferred on expressions, performances, attitudes, and statuses by their

playing a suitable role in a system of discursive normative social practices. The idea of normative statuses as *instituted* by practical attitudes, which has been put on the table in this chapter, should be distinguished from the idea of their intentional contents as *conferred* by the social practices in which those statuses and attitudes play a role. As the terms are used here, the institution of status by attitude has to do solely with pragmatics, the study of the practices in which discursive norms are implicit. The conferral of content by practice has to do with the relation between such pragmatics and semantics, which is the study of conceptual contents.<sup>89</sup> The raw materials for a pragmatics that have been assembled here are employed, in Chapter 3, to contribute to both explanatory projects. The next chapter accordingly begins the investigation of concept use and intentional contentfulness.

### Appendix: Wittgenstein's Use of *Regel*

It should be admitted that Wittgenstein's own terminology in some ways obscures the very point he is after in the regress-of-rules argument. For he uses "rule" extremely broadly, to cover much more than is allowed in the usage endorsed here. According to this latter usage, rules are discursively articulated and propositionally contentful; they determine what is correct by describing the correct performances, *saying* what must be true of a performance for it to be correct. By contrast, Wittgenstein uses "rule" (*Regel*) in at least three importantly distinct senses. First is the sense that coincides with the usage preferred here: rules explicitly say what one is to do and are consulted as such by those who follow them—the rule followers' performance is governed by their understanding of the concepts used to characterize what they ought, according to the rule, to do. Second, he sometimes uses "rule" to describe *whatever* guides or is consulted by those whose behavior is being assessed, whether or not it is discursively or conceptually articulated. Finally, he even sometimes talks about following a rule whenever someone's behavior is subject to normative assessment, whenever responsibility to proprieties of conduct is attributed, regardless of whether there is anything the one "following the rule" is aware of or consulting, or being guided by in determining what to do.

In one central text, Wittgenstein offers two senses in which games may be said to be played according to rules: "The rule may be an aid in teaching the game . . . Or it is an instrument of the game itself—Or a rule is employed neither in the teaching nor in the game itself; nor is it set down in a list of rules. One learns the game by watching how others play. But we say that it is played according to such-and-such rules because an observer can read these rules off from the practice of the game—like a natural law governing the play."<sup>90</sup> The first is following a rule in the sense in which that phrase is used



here, according to which one must understand what the rule says and then try to produce performances that the concepts it employs properly apply to. The other, corresponding to the third of the senses distinguished above, is totally external, involving norms that are only in the eye of the beholder, as the remark about natural laws indicates. These are the two senses that Kant distinguishes as acting according to a *conception* of a rule, as agents do, and merely acting according to rules, as inanimate objects do.

In the context of the regress-of-rules argument, this third sense of rule-following, in which it coincides with simple regularity, must be marginal—the question of how to understand a way of grasping a rule that is not an interpretation hardly arises for inanimate objects as they act according to the laws of physics. This is the sense that seems to be involved in the discerning of rules wherever it would be correct to apply ‘same’ or ‘agreement’.<sup>91</sup> Wittgenstein is even willing to appeal to this sense in such outré (according to the usage preferred here) cases of “rule-following” as those involving rules relating pain to pain-expressing behavior.<sup>92</sup> These would seem to be cases in which the rule is entirely in the eye of the beholder, who takes there to be a regularity. Insofar as they are not, these are cases of the second sort, where the performer is being guided by something, but not by something explicit and articulate. It is in this sense that he is willing to call tables of colors and even signposts “expressions of rules.”<sup>93</sup> He seems to call it “following a rule” wherever there is some object whose features it would be appropriate to cite in justifying one’s performance, exhibiting it as appropriate or correct. Though in one place he seems to be careful not to call a map a rule,<sup>94</sup> in others he is even willing to say a line can function as a rule,<sup>95</sup> and a line is clearly not a saying of any sort.

This multiplicity of senses cries out for the distinctions to be explicitly marked terminologically, which is the intent of the relatively more regimented uses of ‘rule’ and ‘practice’ that is employed here. The important point is that there is a way of grasping a rule that is not an interpretation. This should be talked about, as Wittgenstein at least sometimes does, in terms of *practices*—grasping a rule without interpreting it is grasping it in practice, rather than by substituting one expression of a rule for another. Most cases of understanding explicit claims and obeying explicit orders should be understood in this way. Such application of a rule is something that can be done correctly or incorrectly. Practices in this sense are the primitive sort of acknowledgment that performances are governed by norms. But according to this way of regimenting the idiom, not all practices are graspings of rules. There are practices that involve the acknowledgment of norms without involving rules at all, except in the sense that others, looking on, may be able to state rules—whose expressions are not available to the practitioners.

According to this way of using the term, rules are explicit statements that specify what is correct and incorrect by *saying* (describing) it. Obeying an

order (one of the activities Wittgenstein most often links with rule-following in the *Investigations*) thus counts as a kind of rule-following. But one ought not to say that there are rules involved at all in any practice that does not involve claiming, judging, and describing, though of course there are proprieties of practice in more primitive 'games'. That 'practice' is not to be restricted to "ways of grasping rules that are not interpretations" for Wittgenstein seems to be clear from one of the passages cited above: "To obey a rule, to make a report, to give an order, to play a game of chess, are customs (uses, institutions)."<sup>96</sup> Here rule-following is explicitly just one example. Making a report is not following a rule, though it is governed by proprieties of practice, nor, typically, is giving an order.

## Toward an Inferential Semantics

---

Sure, he that made us with such large discourse,  
Looking before and after, gave us not  
That capability and god-like reason  
To fust in us unused.

SHAKESPEARE, Hamlet

---

### I. CONTENT AND REPRESENTATION

#### 1. *Intentionality: Propositional and Object-Representing Contentfulness*

Taking or treating someone as one of us may be called *recognizing* that individual.<sup>1</sup> According to the construal of recognition being developed here, taking or treating as one of us is adopting a certain kind of intentional stance. It requires first of all interpreting the one it is directed at as the subject of intentional states. But taking someone as one of us also requires, it was suggested, interpreting that individual as an intentional interpreter—as able to *attribute* intentional states, and so as able to adopt toward others just the same sort of attitudes out of which that very stance is constructed. The previous chapter assembled some raw materials for an account of the normative significance of the intentional states we attribute to each other—and take each other to attribute to each other—in adopting the attitudes of mutual recognition that institute the status of community membership, of being one of us.

Before such an account is presented, in the next chapter, it is necessary to look more closely at the sort of *content* that sets apart—as distinctively *intentional*—the states and statuses (and therefore the attitudes) that are attributed when we recognize someone. For intentional states are intention-

ally contentful states, and the theoretical job of the contents they are taken to have is precisely to determine, in context, the particular significance of being in or attributing the states those contents are associated with. As the terms are used here, *semantics* is the study of such contents, and *pragmatics* is the study of the force or significance of the states, attitudes, and performances that have those contents. Accordingly, to fill in the details of a story about the normative character of the *pragmatic significance* of intentionally contentful states, attitudes, and performances, an inquiry into the nature of their *semantic contents* is called for.

Brentano, who brought the term 'intentionality' back into modern usage, defines it this way: "Every mental phenomenon is characterized by what the scholastics of the Middle Ages called the intentional (also mental) inexistence of an object, and what we could call, although not in entirely unambiguous terms, the reference to a content, a direction upon an object (by which we are not to understand a reality in this case), or an immanent objectivity. Each one includes something as object within itself, although not always in the same way. In presentation, something is presented, in judgment something is affirmed or denied, in love loved, in hate hated, in desire desired, etc."<sup>2</sup> "Not always in the same way" indeed. "Intentional object" as used here involves assimilations along two dimensions. First is a dimension most clearly picked out in terms of grammatical categories: what is affirmed or denied in a judgment is something expressible by the use of a declarative sentence, while what is loved or hated is something referred to by the use of a singular term. The former may be called 'propositional contentfulness'. It is typically expressed by the use of a declarative sentence and is ascribed by the use of a 'that' clause appended to a specification of the contentful state or attitude, as in "the belief *that* Carlyle wrote *Sartor Resartus*" or "desiring *that* Pufendorf's reputation be rehabilitated." The latter may be called 'object-representing contentfulness'. It is typically expressed implicitly by the use of a singular term as a grammatical direct or indirect object, and it is attributed explicitly by using terms such as 'of' or 'about', as in "a belief *about* Carlyle" or "desiring something *of* Pufendorf's reputation" (for example that it be rehabilitated).

Putting these two sorts of contentfulness in a box together is not just an idiosyncrasy of Brentano's. Searle, for instance, offers this pretheoretical delineation of the subject matter of his book *Intentionality*: "If a state *S* is Intentional then there must be an answer to such questions as: What is *S* about? What is *S* of? What is it an *S* that?"<sup>3</sup> To insist on distinguishing these sorts of content in the way indicated above is not yet to diagnose a confusion in remarks like this. There is no confusion insofar as propositional and object-representing contentfulness ought to be understood as species of a genus.

An approach to the characterization of that genus is not far to seek. Stalnaker speaks for the dominant tradition in offering this formulation:

"The problem of intentionality is a problem about the nature of representation. Some things in the world—for example pictures, names, maps, utterances, certain mental states—represent or stand for, or are about other things—for example people, towns, states of affairs."<sup>4</sup> The genus, it is suggested, is *representational* content. Indeed, Stalnaker, like others, is comfortable talking interchangeably about "intentional or representational states." This basic insight should be accepted to this extent: it is clear that intentionality has a representational dimension and that to understand intentional contentfulness one must understand representation.

A common response to this insight is to envisage an explanatory strategy that starts with an understanding of representation and on that basis explains the practical proprieties that govern language use and rational action. It is not clear, however, that a suitable notion of representation can be made available in advance of thinking about the correct use of linguistic expressions and the role of intentional states in making behavior intelligible. The temptation to think otherwise is connived at by insufficient appreciation of some of the fundamental criteria of adequacy to which an account of the representational dimension of intentional contentfulness must answer. It is important to keep in mind the explanatory challenges faced by a semantic theory that appeals to representation as its basic concept, and some of the ways in which those explanatory obligations are liable to be unobtrusively shirked. To point these out is not to show that they cannot be satisfied—that representational explanatory strategies are in principle broken-backed. It is merely to guard against the danger that such an explanatory starting point may recommend itself in virtue of its apparent immunity to difficulties it has not squarely confronted.

A particularly unhelpful way of pursuing the representational semantic explanatory strategy is to model representation on *designation*. The designational model is objectionable on two grounds connected with the distinction of grammatical category between sentences and subsentential expressions such as singular terms. First, it assumes that the relation between a singular term and the object it picks out or refers to, for instance that between a name and its bearer, is antecedently intelligible—that the notion of tagging or labeling something can be made sense of before one considers the use of such tags or labels in *saying* something (paradigmatically, in making a claim). In this way, the strategy runs afoul of the principle of the pragmatic priority of the propositional, which is discussed further along.

Second, it assumes that the notion of representation as reference picked out in this way for the category of singular terms and predicates can be univocally and unproblematically extended to apply to the category of sentences. Sentences are understood as representing states of affairs, in the same sense that singular terms represent objects (and in the same sense that predicates represent properties or sets of objects). The notion of representation, conceived as designation, is then supposed to make the grammati-

cal distinction between singular terms and sentences intelligible by appealing to the ontological distinction between objects and states of affairs. Even if it is granted that there is a clear sense in which singular terms such as names and marks on maps represent particular objects, for instance individual people and cities, it does not follow that it is possible to introduce the category of states of affairs as what is in the same sense represented by declarative sentences and 'that' clauses. Nor ought it to be assumed that the ontological category of states of affairs can be made intelligible apart from and in advance of explaining the use of declarative sentences and the 'that' clauses used to report such uses in indirect discourse.

## 2. Two Senses of 'Represents'

Introducing the notion of states of affairs as the kind of thing represented by declarative sentences requires sensitivity to the second dimension of assimilation involved in Brentano's idiom. For one must be careful not to confuse what is *represented* by sentences with what is *expressed* by them. This is a familiar point, but it is worth emphasizing. As Brentano acknowledges by appending to his phrase "direction upon an object" the qualification "by which we are not to understand a reality in this case," 'represent' is ambiguous between two intimately related but importantly distinct senses. Searle puts the point this way: "'About' . . . has both an extensional and an intensional-with-an-s reading. In one sense (the intensional-with-an-s), the statement or belief that the King of France is bald is about the King of France, but in that sense it does not follow that there is some object which they are about. In another sense (the extensional) there is no object which they are about because there is no King of France. On my account it is crucial to distinguish between the content of a belief (i.e. a proposition) and the objects of belief (i.e. ordinary objects)."<sup>5</sup> Thus as Searle sets things up, for a statement or belief to have content is for it to represent or be about something in the 'intensional' sense, while for it to have an object or objects is for it to represent something in the 'extensional' sense.<sup>6</sup> The relation between the two senses emerges more clearly if one or the other is taken as primary and the remaining one specified in terms of it. Thus if 'represent' is reserved for the sense in which one can represent only what in fact exists, whether it be in the category of objects or of states of affairs—actual objects corresponding to singular terms and actual states of affairs corresponding to true claims—then the other sense can be picked out as *purporting* to represent. The other way to do things would be to use 'represent' even in the cases where nothing exists to be represented, where there need be no object or state of affairs as represented for there nonetheless to be a representing. When something does exist as represented, the representation might be called *successful* or correct.

An account of contentfulness in terms of representation needs to explain

both of these senses in which something can be a representing, and it needs to explain their relation to one another. It is clear that if contentfulness in general is to be identified with representational contentfulness, that is, with being a representing, then 'representing' should be understood as purported representing in a sense that contrasts with successful representing. For it makes sense to wonder whether, or to believe that, there is a present king of France or Schelling was the greatest German philosopher, even if it turns out that no object or state of affairs corresponds to that contentful state. A theoretical idiom that shrinks the scope of purported representing until it coincides with that of successful representing has no room for the notion of error, of representation that is incorrect or mistaken; and a notion of representation so thin as to preclude assessments of correctness provides no basis for any recognizable concept of intentional content.

A theoretical idiom that, on the contrary, expands the scope of successful representing until it coincides with that of purported representing is equally unpromising. The result of holding purported representing fixed and failing to distinguish successful representing from it is Meinongianism—commitment to a vast realm of entities, most of which do not exist, including many that *could* not exist. The trouble with taking it that there is something that is successfully represented by every purported representing is not just that it involves commitment to a luxuriant ontology; ontological self-indulgence is a comparatively harmless vice. But it can be symptomatic of a failure to shoulder an explanatory burden. In this case it evidently (and ultimately unhelpfully) transforms the demand for an account of the relation between correct and incorrect, unfulfilled or merely purported and actually successful representing, into a demand for an account of the relation between the statuses of what is represented in the two cases: between mere subsistence and robust existence. Ontological postulation can no more provide an explanation by itself in this case than it could in the one just considered, where the issue was an account of the relation between the sense in which singular terms are representationally contentful and the sense in which sentences are. (Of course, no more in this case than in that one does a commitment to taking the representational dimension of intentional content seriously entail going on to make such a mistake; it is important to recognize the temptation in order to resist it.)

Brentano, who did not make the mistake of his student Meinong, indicates some of the difficulties faced by such an attempt to ontologize the distinction between correct and incorrect representation by holding to a univocal sense of 'represents' and construing the distinction as a difference between two different sorts of representable: "It would be paradoxical to the highest degree to suppose that you could promise to marry an *ens rationis* and then keep the promise by marrying an actual, concrete particular."<sup>7</sup> It is disastrous to put the notion of successful representing in place of that of purporting to represent, that is, to have it play the role of necessary condition for content-

fulness. But while the two senses of 'represent' or 'about' must not be run together (from either direction), there is also reason not to want them to be driven too far apart. Purporting to represent is intelligible only as purporting to represent *successfully* or *correctly*. If what would make the representings successful has no part to play in determining the purport or content of those representings, it is hard to see how assessments of correctness could even get a grip on them. The trouble then is not just that of skepticism about justification, in the Cartesian mode. If all our ideas could have just the content-as-representational-purport that they do, even though the rest of the world, the representeds those ideas purport to represent successfully, were entirely different from what it is represented (purported) to be, how could we ever be justified in taking ourselves to be correct? The difficulty that looms is more serious still, threatening not just the cogency but even the comprehensibility of the picture of states and attitudes as contentful in virtue of their representing or being about the way things are. For the very notion of representings so much as purporting to be about representeds becomes unintelligible.

Acknowledging this distinction between representational purport and representational success is one of the theoretical jobs Frege assigns to his paired semantic concepts *Sinn* and *Bedeutung*. A sign is contentful insofar as it *expresses* a sense. A thought is the sense, the propositional content expressed by a declarative sentence. To say that it is true or false—to assess it along the dimension of correctness semantically relevant to thoughts—is to classify it in terms of the result of applying a function to objects serving as the arguments of that function, where both the function and the objects are picked out as those referred to by components of that sense. The structure of the later Frege's semantic project accordingly encompasses accounts both of what it is to express a sense and of what it is for that sense to be *correct* in terms of how things are with what it represents. An utterance or inscription *expresses* a sense, for example a thought, and it is the sense expressed that then *refers* to objects, the thought that represents them as instantiating properties and standing in relations.<sup>8</sup> This idiom avoids the dangerous ambiguity inherent in talking about propositions as *represented* by sentences. For that way of talking is liable to be misunderstood as involving the identification of propositions with the facts or states of affairs successfully represented by true claims (according to the representational model of contentfulness) rather than with the claims or purported representations expressed by sentences.

### 3. Representational Uptake

The notion of representational *purport* implicitly involves a notion of representational *uptake* on the part of some consumer or target of the purporting. It is only insofar as something can be *taken* to be a representation



that it can purport to be one. For purporting to be something is putting oneself forward as aptly or appropriately taken to be that. The purport is veridical or spurious (for instance the representation is successful or misleading) accordingly as the taking it invites is correct or incorrect. That grasp of something as a representation is coordinate with representational purport is the point Dennett is making when he says: "Something is a representation only for or to someone; any representation or system of representations thus requires at least one user or interpreter of the representation"<sup>9</sup> (using "representation" to mean purported, not necessarily successful, representation). It was pointed out in the previous paragraph that according to the representational model of contentfulness being considered, representational purport is what is *expressed* by a representing, for instance a sign design, rather than what is *represented* by it if it is successful. The present point is then that talk of what is expressed is intelligible only in the context of talk of the activity of grasping what is expressed. By widening the focus a bit, this can be seen to be the manifestation (within the representational construal of contentfulness) of the general point that *meaning* and *understanding* are coordinate concepts. The notion of representational purport is one way of rendering what must be understood in grasping the content of an intentional state, attitude, or performance.<sup>10</sup> Representational purport and the understanding that is its uptake must both be explained in order to make an account of intentional contentfulness in terms of representation work. As Kant says: "The understanding, as a faculty of knowledge which is meant to refer to objects, requires quite as much an explanation as the possibility of such a reference."<sup>11</sup>

Looking back from the vantage point won for us by the later Wittgenstein, it is possible to see that one of the unfortunate emphases that Descartes imposed on the representationalist tradition is the privileging of *knowledge*, and therefore *successful* representation, as a topic of inquiry, over *understanding*, and therefore *purported* representation. For Descartes, representational purport, being "as if of" something, is an intrinsic and characteristic property of *pensées* (that is, specifically mental acts). He does not offer an account of what it is for a mind to grasp such purport, for it to take or treat an idea as being of or about something. He is concerned with how one might become entitled to a commitment to something that has objective (in his, neo-Scholastic sense) reality in one's thoughts having also formal reality outside them. He is not concerned with what the mind's taking one thing or sort of thing rather than another (or rather than nothing at all) as having objective reality in one's thoughts itself consists in. Representational purport, "the objective reality of things in thoughts," and its corresponding uptake by the mind whose thoughts they are serve Descartes as unexplained explainers. So the content of the representational commitments to which the mind's entitlement is at issue is never clarified. A representational model of contentfulness cannot rest with an account of

successful representation—not even if it is accompanied by a vindication of the right to believe that purported representation is often or even generally successful. It requires also an account of representational purport, and that requires an account of the uptake, grasp, or understanding of such purport.

It would of course be a blunder, of a familiar kind, to understand that uptake in general as consisting in *interpreting* something *as* a representation, in Wittgenstein's sense of 'interpreting'. Taking something as a representation must not be parsed in terms of the adoption of explicitly contentful attitudes or intentional states such as belief. If being a consumer of representational purport, taking something as a representation of something, is understood as believing of it that it correctly represents (or equally if the purport is understood as intending that it do so), then an infinite explanatory regress is generated by the possibility of querying the nature of the representational purport ('that . . .') and success ('of . . .') such a belief exhibits. There must be some way of understanding something as a representation that consists not in interpreting it (in terms of something else understood as a representation) but in taking, treating, or using it in practice *as* a representation. To understand what it is for red dots on a map to purport to represent cities and wavy blue lines to purport to represent rivers, the theorist must look to the practice of using a map to navigate. If such purport is to provide a model applicable to representational purport in general, that practice must admit of construals that do not appeal to the formation of propositionally contentful beliefs. The practice must be intelligible in terms of what counts as following it or going against it in what one actually does: the way it guides the behavior of those who can use maps.

The absence of a nonregressive account of what it is to take, treat, or use something *as* a representation of something else is the source of another traditional sort of dissatisfaction with the representationalist paradigm of contentfulness. It lies behind Rebecca West's irritated response to the "mind as the mirror of nature" model that it is hard to see why one would want a copy of the universe: "One of the damn things is enough." Progress in understanding intentional contentfulness is made by invoking representational relations only in the context of an explanation of what it is that makes representings graspable or intelligible as representings in a way in which what is represented is not. That is a matter of the uptake or consumption of representational contentfulness. Apart from the representational purport it expresses, and which is there to be grasped, a representing is just another bit of worldly furniture, like what it represents. Why is not confronting a map as well as terrain just adding one more thing to be baffled about? Invoking a relation (for instance some sort of isomorphism) between representing and represented does not by itself contribute to the task of explaining what the intelligibility of the representing consists in—why one of the damn things is not enough.

#### 4. *Expression and Representation*

Restricting attention for the moment to the propositional contents characteristic of intentional states such as belief, it has been suggested that it is no use asking what a proposition (or propositional content) is, without asking what it is for a sentence to *express* a proposition, or for a state to have one as its content. Just so it is no use asking what it is for a proposition to be *true*, or a representation to be *successful* or correct, without asking what it is to express one—what *purporting* to represent consists in. And it is no use asking what it is to *express* a proposition or other content (to *purport* to represent), without asking what it is to *grasp* or *understand* such purport. An account is required not only of how representings are distinguished from and related to representeds (in successful representing). An account is needed also of the representational content representings express—their representational purport. And that requires an account of the *attitude* of grasping such purport: of taking, treating, or using a representing as a representing, of acknowledging or attributing to it in practice its representational purport.

The treatment of representational content in upcoming chapters centers on an account of this practical attitude. Becoming entitled to use a concept of intentional content involves a twofold explanatory task: to say what it is to *express* a *propositional* content in general, and then to say what more is required specifically for the content expressed to *represent* something *objective*, in the way that matters for empirical science. Furthermore, each of these must include an account of what those who exhibit and attribute states, attitudes, and performances with such contents must do in order to count as taking or treating them in practice as contentful in those ways.

This is a request that can sensibly be addressed to Wittgenstein, as well. Even his sustained, penetrating discussions do not offer an account of what distinguishes language games within which states and performances acquire specifically *propositional* significances (the only ones that, by the lights of this work, deserve the title 'Sprachspiel'), nor of what distinguishes those within which states and performances acquire specifically *representational* significances. He argues against understanding the contents determining the significances of all states and performances in terms of representational content. For one ought not simply to presume on syntactic grounds that terms are used to refer (or fail to refer), predicates are used to describe or characterize (or misdescribe), and sentences are used to claim (truly or falsely). Instead of asking what object is being referred to by the term, what property is being ascribed by the predicate, and what fact would make the sentence true, one ought first to look at the *use* of the expressions, to see if the putative referings, characterizings, and claimings in fact play a practical role that is best understood in terms of such contents. For many expressions

that might have been thought to be doing the jobs just mentioned ('sensation', 'intending', 'beetle', 'must', 'true', 'I am in pain', . . .), consideration of their use shows that another account of the content of the putative referings, characterizings, and claimings is more appropriate.

The idea that not all contentful expressions play a straightforwardly representational role is a development of a line of thought that is already important in the *Tractatus*. Some previous varieties of logical atomism had distinguished themselves by their insistence that the only way any expression, sentential or not, could have content or contribute to the content of an expression of which it is a part is by standing for or representing something.<sup>12</sup> Thus, not only did these views grasp the nettle of commitment to negative and conditional facts, they also were committed to 'not' and 'if . . . then . . .' standing for some element in a complex state of affairs. The undertakers of such commitments are admirable more for their conceptual heroism than for their good sense.

In the *Tractatus*, Wittgenstein showed that one could best treat logically compound propositional contents as representing states of affairs by not treating every contentful expression (every one whose occurrence is significant for determining the state of affairs represented by the whole) as itself having its content in a representational way, by standing for something. Purely formal vocabulary, paradigmatically logical vocabulary, is contentful but does not itself stand for anything. (Kant and Frege had of course earlier shown the possibility of this sort of approach.) The opening sections of the *Investigations* argue along just these lines: not every piece of a representation contributes to its content by itself representing, and not every move in a language game is a representing of something. But Wittgenstein does not explain what one must do to be using an expression to refer, characterize, or claim (the features of use he associates with representational content), nor does he explain what is required for something caught up in a language game to express a specifically propositional content.

The notion of expression—of making propositionally explicit—shows up at two different levels in what follows. First, one who adopts the pragmatist's approach to intentionality owes an account of the practices that ultimately confer explicitly *propositional* content on the states, attitudes, and performances that play appropriate roles in those practices. This is an account of the implicitly normative practices in virtue of which anything at all can be made explicit as the content of a possible claim or belief *that p*. Such a theory should explain what it is for a performance, paradigmatically but not exclusively the tokening (by speaking or writing) of some linguistic item, to express an intentional content. And it should explain the relation between such expressions and the possession of content by states or attitudes somehow related to them. That is, it should explain what it is to express or exhibit a specifically *propositional* content—intuitively, one that could be true or false. Furthermore, it should explain what it is to express or exhibit a content

that purports to *represent* something, and it should explain the relation between representing states of affairs and representing particular objects. As has been pointed out, a necessary part of explaining the expression of contents with representational purport is explaining the *grasping* of such contents, the uptake that is the other side of such purport.

The second level concerns not the making of ordinary claims but the formulation of rules or principles. The regress-of-interpretations argument shows that the intellectualist tradition erred in treating the explicit form of norms as fundamental. But once a notion of propositional explicitness has been brought onboard in terms respectable according to pragmatist scruples, the fact that the contentful norms implicit in practical doings can be expressed in rules, claims, and interpretations that say or state explicitly what is implicit in those practical proprieties itself still stands in need of explanation. An account is needed of what it is to make explicit in the form of something that can be *said* or *thought* what is otherwise merely implicit in what is *done*. At this level, the implicit proprieties of practice that make it possible to make propositionally explicit claims are themselves made propositionally explicit in the form of rules or principles. A theory of *expression* accordingly is to explain how what is *explicit* arises out of what is *implicit*. In the first instance, it must explain how propositional content (the form of the explicit) is conferred by norms that are implicit in discursive practice—that is, what proprieties of use having such a content consist in. Then it must show how those same implicit, content-conferring norms can themselves be made explicit in the form of rules or principles.

### 5. From Practice to Content

These two challenges are addressed in the rest of the work. First, what role must states, attitudes, and performances play in (as it turns out, social) practice for it to be correct to interpret them as being *propositionally* contentful? That is, how are propositional contents conferred by practice? What proprieties of practical employment does possession of such content consist in? As already suggested, any answer must specify what it is for the practitioners themselves practically to take or treat states, attitudes, and performances of others and of their own as having such contents, and thereby to confer those contents on them. Chapters 3 and 4 develop a response to these questions. Second, what must be true of such contentful states, attitudes, and performances for it to be correct to interpret them as *representing* objects and objective states of affairs? Again the answer must specify what it is for the practitioners themselves in practice to take or treat those states, attitudes, and performances as having such contents, and so by their practice to confer such contents on them. Part 2—particularly in Chapter 6 (on the representation of objects by singular terms) and Chapter 8 (on objective

representation)—presents an account of these phenomena, within the framework introduced in Chapter 3.

The practical uptake of specifically representational purport must include normative *assessment* of states, performances, and expressions—assessment of their specifically representational correctness. (Of course, on pain of the familiar regress, such assessment must not be understood as in every case consisting in judging *that* a representation is correct; besides such propositionally explicit attitudes there must be practically implicit ones.) Treating something as a representation involves acknowledging the possibility that it *misrepresents*—that the representational taking is a *mistaking* (the object represented does not exist, the state of affairs represented does not obtain). It is these attitudes of distinguishing in practice between representations that are taken to be correct and those taken to be incorrect that forge the connection between the notions of representational purport and representational success.

Practical representational uptake of representings—treating objects, states, or performances as purporting to be correct representations of objects and facts—consists in taking them to be takings: taking them to express attitudes concerning what there is and how things are. That they are accordingly essentially liable for assessment as to their representational success (that they in a characteristic way answer to how things actually are for their correctness) means that such uptake incorporates an implicit distinction between representational *attitude* (how things are taken to be by what is treated as a representation) and representational *status* (how things actually are, which determines the success or correctness of that attitude). Thus the normative pragmatic distinction between status and attitude is central to the intelligibility of fundamental semantic concepts. It is reflected in the distinction between representational purport and representational success.

The *objectivity* of representational content is a feature of the practices of assessing the correctness of representations. The status of representings as correct or incorrect, successful or unsuccessful, depends on how things are with what is represented, rather than on the attitudes of representers. What is distinctive of specifically *representational* correctness is this objectivity—the way in which assessments of representational correctness take representings to answer to what is represented, rather than to how what is represented is *taken* to be. It is the way in which the status being assessed outruns any particular attitude toward it. Understanding the objectivity of representational content requires understanding this particular structure of authority and its acknowledgment—what it is for those assessing the correctness of representings to cede authority over them to what is represented, to treat their correctness in practice as determined by those represented. Again, one lesson is that the representational dimension of semantic content cannot be understood apart from the normative pragmatic context in which it is embedded and in which it is accorded its characteristic significance.

It should be clear that the remarks in this section are not meant to have the force of arguments against treating representation as a central semantic category. Rather, they present some general criteria of adequacy for an account of this important semantic notion. By doing so, however, they do offer reasons not to treat representation as a semantic *primitive*, as an unexplained explainer. The next section shows why the role of semantic concepts in pragmatics (the proper use of language and the appropriate role of intentional states in rational action) dictates approaching semantics in the first instance through the notion of *propositional* contentfulness. The rest of the chapter then motivates an approach to propositional contentfulness that begins with the *inferential* articulation of the social practice of giving and asking for reasons. The following chapter presents a particular model of those social practices (in terms of deontic scorekeeping) and shows how they can be understood as at once instituting discursive commitments and conferring propositional contents on them. In Part 2, that framework is extended to include *representational* content, both of the sort expressed by sentences and that expressed by subsentential expressions. It concludes with a discussion of the social and inferential articulation of discursive practice in virtue of which the contents it confers are properly understood as involving an *objective* representational dimension.

## II. THE PRIORITY OF THE PROPOSITIONAL

### 1. *Kant on Judgment as the Form of Awareness*

It is appropriate to begin by addressing propositional contents because of what can be called the *pragmatic priority of the propositional*. The pre-Kantian tradition took it for granted that the proper order of semantic explanation begins with a doctrine of *concepts* or *terms*, divided into singular and general, whose meaningfulness can be grasped independently of and prior to the meaningfulness of judgments. Appealing to this basic level of interpretation, a doctrine of *judgments* then explains the combination of concepts into judgments, and how the correctness of the resulting judgments depends on what is combined and how. Appealing to this derived interpretation of judgments, a doctrine of *consequences* finally explains the combination of judgments into inferences, and how the correctness of inferences depends on what is combined and how.

Kant rejects this. One of his cardinal innovations is the claim that the fundamental unit of awareness or cognition, the minimum graspable, is the *judgment*. "As all acts of the understanding can be reduced to judgments, the understanding may be defined as the faculty of judging."<sup>13</sup> For him, interpretations of something as classified or classifier make sense only as remarks about its role in judgment. A concept just is a predicate of a possible judgment,<sup>14</sup> which is why "the only use which the understanding can make of

concepts is to form judgments by them."<sup>15</sup> Thus for Kant, any discussion of content must start with the contents of judgments, since anything else only has content insofar as it contributes to the contents of judgments. This is why his transcendental logic can investigate the presuppositions of contentfulness in terms of the categories, that is, the "functions of unity in judgment."<sup>16</sup>

The understanding is the active cognitive faculty, the faculty of spontaneity—understanding is something *we do*. "We have before given various definitions of the understanding, by calling it the spontaneity of knowledge (as opposed to the receptivity of the senses), or the faculty of thinking, or the faculty of concepts or of judgments; all of these explanations, if more closely examined, coming to the same."<sup>17</sup> What we do is synthesize, bring things into a unity—that is, subject them to rules or concepts. What we do, as opposed to what happens to us, is to judge. Although synthesis happens at other levels than that of judgment (there is synthesis in intuition and imagination also), that synthesizing activity is an aspect of judging. "The same function which imparts unity to various representations in one judgment imparts unity likewise to the mere synthesis of various representations in one intuition, which in a general way may be called the pure concept of the understanding. The same understanding, and by the same operations by which in concepts it achieves through analytical unity the logical form of a judgment, introduces also, through the synthetical unity of the manifold in intuition, a transcendental element into its representations."<sup>18</sup> Thus all our cognitive activity consists of judgment and aspects of that activity. Any content that can be discerned in any category is derivative from the content of possible judgments, that is, from propositional content. Kant's *pragmatics*, or theory of cognitive *activity*, determines the fundamental unit of his *semantics*, or theory of the *contents* of cognitions.

## 2. Frege and Wittgenstein

This insight into the fundamental character of judgment and so of judgeable contents is lost sight of by Kant's successors (indeed it could be argued that appreciation of it is still missing from such broadly semantic traditions as semiotics and structuralism). It is next taken up by Frege. Looking back over his lifework in 1919, he picks out this point as basic to his orientation: "What is distinctive about my conception of logic is that I begin by giving pride of place to the content of the word 'true', and then immediately go on to introduce a thought as that to which the question 'Is it true?' is in principle applicable. So I do not begin with concepts and put them together to form a thought or judgment: I come by the parts of a thought by analysis [*Zerfällung*] of the thought."<sup>19</sup> Already in 1870 in the *Begriffsschrift*, Frege introduces "contents of possible judgment" or "judgeable contents" in the second paragraph and subsequently defines other sorts



of contents in terms of them. In an essay explaining the *Begriffsschrift* he summarizes this approach: "I start out from judgments and their contents, and not from concepts . . . instead of putting a judgment together out of an individual as subject and an already previously formed concept as predicate, we do the opposite and arrive at a concept by splitting up the content of a possible judgment."<sup>20</sup> The concept of a function, which stands at the center of Frege's technical contribution to semantics, is introduced in the *Begriffsschrift* as an element in his substitutional methodology for decomposing contents of possible judgment.<sup>21</sup> In the *Grundlagen* Frege continues to follow this Kantian line in insisting that "we ought always to keep before our eyes a complete proposition. Only in a proposition have the words really a meaning . . . It is enough if the proposition taken as a whole has a sense; it is this that confers on its parts also their content."<sup>22</sup> Frege holds this view because of the importance he assigns to the concept of *truth*; to talk about an expression as contentful is to talk about the contribution it makes to the truth-value of thoughts or propositions in which it occurs.

It is sometimes thought that Frege gave up his commitment to the primacy of the propositional by the late 1880s, when he began to assimilate sentences technically to singular terms under the heading *Eigennamen*, which includes everything except functional expressions. Such a view overlooks the very special role that sentences, as 'names' of truth-values, continue to play for him, even in the *Grundgesetze*. The importance of truth, and therefore of thoughts (the contents expressed by declarative sentences), continues to be emphasized at every stage in Frege's development. In his long 1914 essay entitled "Logic in Mathematics," he is still maintaining "that the name should designate something matters to us if and only if we are concerned with truth."<sup>23</sup> This is the same view that he had endorsed in his classic essay "Über Sinn und Bedeutung": "But now why do we want a proper name to have not only a sense, but also a reference [*Bedeutung*]? Why is the thought not enough for us? Because, and to the extent that, we are concerned with its truth value . . . It is the striving for truth that drives us always to advance from the sense to the reference."<sup>24</sup> In the context of such a view it is clear that the assimilation of sentences to singular terms as both having objects as *Bedeutungen* can in no way undercut the fundamental role played by truth-values, and so by the propositional contents that bear them. In that same essay he says that what is needed for a name to have content (express a sense) is that it "belong to a sufficiently complete totality of signs."<sup>25</sup> Given his views about identity, this means a system of signs that includes sentences in which the name occurs, and also further sentences that result from them by substituting other names for the ones in question. The totality of signs must include sentences, because to have a sense is to purport to have a *Bedeutung*, and as just indicated, such purport arises only in the context of concern with truth, because "anyone who seriously took the sentence to be true or false would ascribe to the name . . . a *Bedeutung*."<sup>26</sup> It is because the

*point* of deploying concepts in thought and talk is to judge, that is, take or treat judgeable contents as true, that such contents are given pride of place in Frege's scheme. As he says in the 1897 fragment on logic: "Every act of cognition is realized in judgments."<sup>27</sup>

Indeed, it can be misleading to focus on the concept of *truth* as what enforces attention to sentences. Frege takes this position because it is only to the utterance of sentences that pragmatic force attaches, and the explanatory purpose of associating semantic content with expressions is to provide a systematic account of such force. "'True' only makes an abortive attempt to indicate the essence of logic, since what logic is really concerned with is not contained in the word 'true' at all but in the assertoric force with which a sentence is uttered . . . the thing that indicates most clearly the essence of logic is the assertoric force with which a sentence is uttered."<sup>28</sup> Talk about the cardinal importance of concern with truth is a dispensable *façon de parler*. What actually matters is the pragmatic attitude of *taking-true* or putting forward *as true*, that is, judging or asserting. Semantic vocabulary is used merely as a convenient way of making explicit what is already implicit in the force or significance that attaches to the content of a speech act or attitude. (An account of just how this explicitation works is offered in Chapter 5, where specifically semantic vocabulary, paradigmatically 'true' and 'refers', is discussed.)

The point that the contents expressed by sentences must play a privileged explanatory role because it is to sentences that pragmatic force attaches has been brought home most forcefully by the later Wittgenstein. The use of sentences is prior in the order of explanation to the use of subsentential expressions because sentences are the only expressions whose utterance "makes a move in the language game." Sentences are expressions whose unembedded utterance performs a speech act such as making a claim, asking a question, or giving a command. That is why even when such a speech act is performed by an utterance that does not manifest the syntactic complexity typical of sentences (a shout of "Rabbit!" or "Fire!" for instance), the utterance should nonetheless be interpreted as a one-word sentence, as meaning what we might express by "Look at the rabbit!" or "There is a fire!"

Referring to something, indicating or naming it, is also something one can do with linguistic expressions; it is a speech act one can perform. But these belong to a class of speech acts that is in an important sense derivative from or parasitic on speech acts involving sentences, paradigmatically claiming, asserting, or putting forward as true. In order to use an expression as a name, to refer to or pick out an object with it, one must be able to use the name to *say* something (paradigmatically, to assert something) about the object referred to, indicated, or named. The significance of taking or treating something as a name, as purporting to refer to an object, consists in how one takes it to be proper to use the expression, and the use of expressions as names is unintelligible except in the context of using expressions containing them as sentences.<sup>29</sup>

### 3. *Semantics Must Answer to Pragmatics*

The primacy of propositional intentional contents also shows up if one considers cases in which the use of language is not to the fore. Intentional interpretation of nonlinguistic organisms—intentional explanation of their behavior by attributing beliefs and desires that make what they do intelligible—also depends on attributing *propositionally* contentful states, attitudes, and performances. Behavior is made intelligible by exhibiting it as rational, given various beliefs and pro-attitudes, and to do that is to exhibit a piece of practical reasoning that is taken somehow to stand behind or be implicit in the behavior. The imputed reasoning shows why an organism with the states or attitudes that provide the premises *ought*, rationally, to behave in the way specified by the conclusion. But *what can serve as a premise in reasoning must have a propositional content*. This point is so important to the present project that the rest of this chapter is devoted to motivating the treatment of this feature, in the next chapter, as a defining characteristic distinctive of the propositional. The intentional interpreter attributes to the cat the belief *that* there is a mouse around the corner from it, and the desire *that* it catch the mouse, and so on. Attributing intentional states so as to render behavior intelligible in the light of them requires attributing propositional contents to them.<sup>30</sup> So propositional contents have a pragmatic priority, not only in the setting of assessments of the significance of speech acts, but also in the setting of attributions of intentional states that do not evidently depend on linguistic practices.

Semantics must answer to pragmatics. The theoretical point of attributing semantic content to intentional states, attitudes, and performances is to determine the pragmatic significance of their occurrence in various contexts. This means settling how linguistic expressions of those contents are properly or correctly used, under what circumstances it is appropriate to acquire states and attitudes with those contents, and how one then ought or is obliged to go on to behave. It is specifically *propositional* contents that determine these pragmatic significances, so it is specifically propositional contents that it is the task of semantic explanatory theories to attribute. Semantic contents corresponding to *subsential* expressions are significant only insofar as they contribute to the determination of the sorts of semantic contents expressed by full sentences. The pragmatic priority of sentence-use to name-use enforces a certain semantic explanatory priority of the contents expressed by sentences to those expressed by names. The task of the next chapter is to develop an account of the practices of using expressions as sentences—paradigmatically to make claims and so to confer specifically propositional contents on those expression uses and on the states and attitudes associated with them (to use them *as* having such contents).

What the theorist associates with states and expressions deserve to count as *semantic* contents only insofar as they play the right sort of role in determining the proprieties of practice governing those states and expres-

sions. It is possible to associate all sorts of abstract objects with strings of symbols in formalized languages, from sets of models to Gödel numbers. Such an association amounts to specifically *semantic* interpretation just insofar as it serves to determine how those strings are correctly used. For example, Tarski's mapping of well-formed formulas of the first-order predicate calculus onto topological domains qualifies as a semantic interpretation of them only because he can derive from it a notion of valid inference, a way of telling what follows from what—that is, a notion of their correct use. Apart from that, it would just be one more algebraic homomorphism.

#### 4. *Two Mistakes the Designational Model Invites*

An account of content in terms of representation must satisfy the requirement that it must show how semantic content so construed matters for the pragmatic significance of what it is associated with. For the reasons indicated above, this demand focuses attention to begin with on the representational rendering of specifically propositional contents. Two difficulties arise at this point: it is not clear how to derive a notion of propositional contentfulness from the designational representational model, and construing content in representational terms requires supplementation by a further story to get to the proper use of contentful expressions and the correct circumstances and consequences of being in contentful states. (In contrast, the explanatory strategy pursued in Chapter 3 begins with an account of the practices within which producing a performance or altering an attitude can have the pragmatic force or significance of making a claim or judgment; the notion of propositional contentfulness is then understood as what is expressed by such acts.)

On the first point, the pre-Kantian representationalist tradition offers no useful account of what is represented by judgments. For this tradition, representational relations hold between *things*. This categorial nominalism of the designational model extends even to predicates, which are understood as 'general names,' standing for universals in the same sense in which singular terms stand for particular objects. Not until Frege's semantic interpretation of predicates as corresponding to functions—and hence as not being names of any sort—would the idea of semantic relations that are not assimilable to the name/named model enter the tradition.

Applied to propositional contents, the hegemony of the designational semantic model results in two characteristic mistakes: assimilating sentences to complex names, and assimilating judging to predicating. Kant provides the raw materials needed to move beyond these conceptions, but even he is not able to free himself entirely from them. That the first is a mistake becomes clear in the context of an attempt to explain the difference between *referring to a complex object*, for instance a squiggly blue line between a round dot and a square one, and *stating a fact* about its components, for instance *saying*

that the squiggly blue line is between a round dot and a square one.<sup>31</sup> This crucial difference can be elided by an incautious assimilation of each to a generic notion of *representing*, for in each case the speaker can be said to be representing something.

This difficulty is merely relocated by the introduction of a notion of state of affairs defined as the sort of thing that is represented by utterances that purport to state facts. Making this move is a version of the attempt to solve ontologically the problem of distinguishing referring from saying or stating; the idea is that each is representing, and the specific differences between them are a matter of the kind of thing represented. At the least, such a strategy demands a careful account of the relation between complex objects and the corresponding states of affairs. Any account along these lines of discourse that purports to state facts by the assertive utterance of declarative sentences is also obliged to tell a story about the states of affairs corresponding to normative claims—for instance to the claim that Kant ought not to have written *Dreams of a Spirit-Seer*, or that anyone committed to the claim that snow is white is committed to the claim that snow is spatially extended.<sup>32</sup>

The second mistake mentioned above as consequent on the unfortunate sway of designational semantic models corresponds to one way in which the representational tradition in semantics has attempted to acknowledge the special role played by the propositional. This strategy depends on a distinction on the side of the activity of representing, rather than (just) the category of thing represented, by distinguishing between representing as referring or naming, on the one hand, and representing as predicating, on the other. The notion that making a claim can be modeled on representing something (particular) *as* something (general), in the linguistic case picking out an object with a singular term and predicating something of it with a general term, has a distinguished history. The next chapter discusses Frege's decisive demonstration that this approach is a mistake. It is not simply a mistake, however. Looking more closely at what is right about this broadly *classificatory* model of consciousness and at how it can be fixed up in response to some fundamental difficulties provides a way into an idiom for talking about semantic content that does not employ representational vocabulary at the outset.

### III. CONCEPTUAL CLASSIFICATION AND INFERENCE

#### 1. *Classification*

An ancient tradition insists that cognition essentially involves generality or universality. Particulars are not directly intelligible as such. Knowing or understanding something particular requires assimilating it to others, taking it to be like them in some way, and so to be an instance of a kind. Kant's account of cognition as beginning with the classification of

intuitions under concepts is a particularly well-developed representative of this tradition. Believing or judging, taking-true in general—for Kant the central sort of cognitive intentional state or act—has as its most basic form subsuming something particular under a universal.<sup>33</sup> In conceiving judgment (the activity of the cognitive subject, the exercise of its faculty of spontaneity, namely understanding) in terms of the *classificatory* employment of concepts, Kant adopts a model that animates as well the thought of the pre-Kantian tradition he inherits—a tradition that had not yet achieved his insight into the privileged role of judgment as the preeminent form of cognitive activity.

That model evidently underlies the epistemologies of both his empiricist and his rationalist predecessors. It forms the common background of their dispute over the source of the universals or concepts by relation to which particulars become intelligible; it is what makes urgent the question whether those universals are formed by abstraction from a more primitive kind of nonconceptual awareness of particulars, or whether on the contrary a grasp of such concepts is a precondition of anything recognizable as awareness at all. Kant follows the rationalists in treating the classificatory account of cognition as a classificatory account of consciousness generally. All awareness is understood as exhibiting the classificatory structure of universal or repeatable concepts subsuming particulars. Where earlier empiricists admit varieties of conscious apprehension short of conceptual comprehension—immediate, nonclassificatory awareness of determinate sense repeatables, for instance—Kant denies apprehension without comprehension, insisting that there must be conceptual classification wherever there is any sort of awareness. Awareness of what is classified and of how things can be classified derives from awareness that consists in classifying.<sup>34</sup>

A pragmatic version of this classificatory model results if it is de-intellectualized, stripped of residual commitments to understanding concepts as *explicit* to the mind—whether in the Kantian form of rules or recipes for it to follow in its synthesizing activity or in the pre-Kantian form as objects of its direct, nonclassificatory awareness. From such a perspective, the roots of conceptual classification are to be found in treating something in practice as being of a certain kind—taking something (particular) *as* something (universal), by behaving toward it in a way that assimilates it to others. Particular objects are classified as belonging together in some respect by being responded to alike in practice. A respect of similarity in what is responded to then corresponds to a repeatable response. Hegel develops such a pragmatic, indeed naturalized, version of Kant's account in the form of an *erotic* theory of the origins of awareness, an account of animal *desire* as the source of classification. As he puts it, an animal classifies some particular as food when it "falls to without further ado and eats it up."<sup>35</sup> Eating something is treating it, responding to it, classifying it in practice *as* food. It exhibits a kind of *practical*, pre-Cartesian awareness of it *as* being of a certain kind. That

repeatable activity on the part of the organism induces a repeatable respect of similarity among the things that tend to elicit that activity.

On this account, classification of particular stimuli as instances of a general kind is *implicit* in what the responding organism *does*. So to each sort of thing that it does, there corresponds a different sort of repeatable proto-concept under which things can be classified: as food, sexual partner, prey, or predator, and so on. That no sort of explicit awareness is presupposed by this sort of implicit practical classificatory awareness or understanding is clear from the fact that all that the latter requires is a reliable differential responsive disposition. For any concrete object displays such dispositions. A chunk of iron reliably responds to some environments by melting, to others by rusting, to still others by falling. In each case it can be understood as classifying that environment, treating it in practice as being of a certain kind, assimilating it to some other possible environments and distinguishing it from others, by responding to it in a certain way.<sup>36</sup>

The Kantian rationalistic strategy of demarcation by *sapience*, awareness, and consciousness in a sense that requires the application of *concepts* would be trivialized by a classificatory model of the use of concepts that indiscriminately discerns classification according to concepts in the responsive regularities exhibited by the antics of every physical system whatsoever. Classification by the exercise of regular differential responsive dispositions may be a necessary condition of concept use, but it is clearly not a sufficient one. Such classification may underlie the use of concepts, but it cannot by itself constitute discursiveness. The chunk of iron is not conceiving its world as wet when it responds by rusting. Why not? What else must be added to responsive classification to get to an activity recognizable as the application of concepts? What else must an organism be able to do, what else must be true of it, for performances that it is differentially disposed to produce responsively to count as applications of concepts to the stimuli that evoke those responses? One dimension of a reply was indicated in the previous chapter—a normative dimension is required, which can underwrite a distinction between correct and incorrect applications of concepts. But many things can be done correctly or incorrectly. The question being asked now is what it is for what is subject to such assessment to be concept use (rather than, say, hammer use, or tooth use).

## 2. Inferential Demarcation of the Conceptual

An easy answer is that the response must be the forming of a belief or the making of a claim, acquiring a state or attitude or producing a performance that has an intentional content. This is of course correct, but unhelpful in the current setting. For the question is precisely what is required for a response to count as contentful in this sense. What is wanted is a characterization that does not appeal to semantic concepts such as *content*

and *concept*. If the issue is put in terms of the semantic concept of representation, it takes the form of inquiring as to what more is needed, beyond being a representation in the responsive-classificatory sense, to be a discursive or intentional representation, one that is conceptually contentful. (According to the idiom being employed, implicit grasp of such contents, of the representational purport they consist in—a grasp to be conceived of as some sort of practical mastery, as a kind of know-how—would then in favored cases count as sapient consciousness or awareness of what is represented as exhibiting a certain character.)

A more concrete way to put the question is to ask, What are the salient differences between a measuring instrument, such as a thermometer or spectrophotometer, and an observer who noninferentially acquires beliefs or makes claims about environing temperatures and colors? Artificial instruments differ from other physical systems, such as chunks of iron, only in having been constructed so that some subset of the partition of possible stimuli into equivalence classes according to the distinguishable responses the instruments are disposed to produce corresponds to some distinction of practical or theoretical significance to the user, who thereby attaches some significance to them. Suppose a spectrophotometer is hooked up to a tape recorder in such a way that it produces a noise of the acoustic type "That's red" when and only when it is irradiated with light of the proper frequency. And suppose that a fanatical human red-reporter nearby has just the same responsive dispositions to produce those noises. That is, the two systems are disposed to respond in the same way to the same stimuli, exhibiting the same noninferential circumstances of application for their responsive classifications of things as red. What makes the noise the one produces merely a signal on the basis of which someone else might conclude that something red is present, while the very same noise, reliably elicited under just the same circumstances from the other, counts as a noninferential report, expressive of a perceptually acquired belief, with an intentional content that includes the concept *red*? To vary the case, suppose the reporter's differential responsive dispositions to call things red are matched by those of a parrot trained to utter the same noises under the same stimulation. What practical capacities of the human distinguish the reporter from the instrument or the parrot? What, besides exercise regular differential responsive dispositions, must one be able to do, in order to count as having or grasping concepts, and so as able to perform not only classification but specifically *conceptual* classification?

Putting things this way makes it clear that what is at issue is a kind of *understanding*. The reporter's response is meaningful—not just, as in the case of the measuring instrument or the parrot, to others, but to the responding reporter personally. The spectrophotometer and the parrot do not understand their responses; those responses mean nothing to them, though they can mean something to us. The reporter understands the response he or she makes, attributes to it a kind of significance that the measuring instrument



and the parrot are oblivious to. The challenge is to explain what sort of practical capacity the relevant kind of understanding consists in, without an ultimately circular appeal to semantic concepts such as intentional content, concept-use, or the uptake of representational purport (treated as an explanatory primitive).

The leading idea of the approach to content and understanding to be developed here is due to Sellars. Sellars's suggestion is that the key element missing from the parrot and the measuring instrument—the difference between merely *responsive* classification and *conceptual* classification—is their mastery of the practices of giving and asking for *reasons*, in which their responses can play a role as *justifying* beliefs and claims. To grasp or understand a concept is, according to Sellars, to have practical mastery over the *inferences* it is involved in—to know, in the practical sense of being able to distinguish, what follows from the applicability of a concept, and what it follows from. The parrot does not treat “That’s red” as incompatible with “That’s green,” nor as following from “That’s scarlet” and entailing “That’s colored.” Insofar as the repeatable response is not, for the parrot, caught up in practical proprieties of inference and justification, and so of the making of further judgments, it is not a conceptual or a cognitive matter at all. What the parrot and the measuring instrument lack is an appreciation of the significance their response has as a reason for making further claims and acquiring further beliefs, its role in justifying some further attitudes and performances and ruling out others. Concepts are essentially inferentially articulated. Grasping them in practice is knowing one’s way around the proprieties of inference and incompatibility they are caught up in. What makes a classification deserve to be called *conceptual* classification is its *inferential* role. It is practical mastery of the inferential involvements of a response, the responder’s understanding it in this sense, that makes the response an intentional state or performance—one having a content for the one whose state or performance it is, and not merely for those using it as an indicator.

### 3. *Holistic Consequences of Inferential Approach to Concepts*

One immediate consequence of such an inferential demarcation of the conceptual is that one must have many concepts in order to have any. For grasping a concept involves mastering the proprieties of inferential moves that connect it to many other concepts: those whose applicability follows from the applicability of the concept in question, those from whose applicability the applicability of the target concept follows, those whose applicability precludes or is precluded by it. One cannot have just one concept. This holism about concepts contrasts with the atomism that would result if one identified concepts with differential responsive dispositions. The capacity to

treat some things as food by eating them need have no particular connection to the capacity to treat other things as dangerous by fleeing them. To treat states or performances as intentionally contentful in the sense of being conceptually articulated involves treating them as situated in a web of proprieties of inferential transitions from one content to another. Knowing one's way around the bit of the web centered on one conceptual content, being able to tell in practice which moves to it and from it are permitted or required and which forbidden, accordingly requires mastery of the proprieties of inference that govern the use of other concepts and contents as well.

By contrast, there is *prima facie* no reason why the fact that some object or property is represented by one simple idea, term, or predicate should be relevant to what is represented by others. Representational relations between nonintentional objects or properties and the intentional representings of them might be treated (as the empiricists in fact treat them) as separate building blocks that, when properly put together, determine what inferences are good in the sense of preserving accuracy of representation. Serving this role seems compatible with these representational relations being quite independent of one another. Knowing what one state or expression represents need convey no information at all about what anything else might represent.

But the inferential notion of semantic content is essentially holistic. Inferences involve both premises and conclusions. The inferential role of one of the premises essentially depends on that of the conclusions, and vice versa. One could not know something about the inferential role of one content without knowing at least something about the inferential roles of others that could be inferred from it, or from which it could be inferred. Contents understood in terms of inferential roles are evidently interdefined in a way in which contents understood in terms of representational purport need not be.

In his masterwork, "Empiricism and the Philosophy of Mind," Sellars exploits this consequence of his insight into the significance of inferential connections to concept-use, even in cases of responsive classification. He argues there that noninferential reports, by which perceptual states are made explicit, cannot constitute an autonomous fragment of a language—one that might be understood though no others are. Observation reports do indeed have a certain priority in the order of *justification* of empirical claims. But they cannot be accorded a similar priority in the order of *understanding* of those claims. Since knowledge requires not only justification but grasp or understanding of the content being justified, there can be no observational knowledge without inference. There can be no purely observational language or set of concepts with respect to which one could then ask whether the decision to append an inferential superstructure is a rational or justifiable one. The rock on which foundationalism founders is accordingly its incapacity to explain what it is to *understand* the significances of elements in the observational justificatory basis. For in order to be able to apply one concept

noninferentially, exercising a disposition to respond differentially to nonlinguistic stimuli, one must be able to apply others inferentially. Unless the response has such an inferential significance, it is not a conceptually contentful response. So the idea of an autonomous language game (or set of practices of applying concepts) consisting entirely of noninferential reports (even of purely mental happenings) is a radical mistake.

The argument does not rule out the possibility of languages or conceptual schemes that are devoid of *theoretical* claims and concepts—that is, that lack concepts that are applicable only as conclusions of inferences. One can have a scheme in which all the concepts have reporting uses and so are in this sense concepts of observables. But they must also have inferential uses. *Red* can be applied either noninferentially, as a response the reporter has been trained to make to a certain kind of visual stimulus, or inferentially, on the basis of entitlement to a prior application of the concept *scarlet*. The conclusion that there can be no conceptually articulated observation apart from inferential capacities holds equally whether what is being reported consists of external observable situations or internal, purely mental happenings. It is this argument that lies at the base of Sellars's critique of broadly Cartesian philosophies of mind.

#### 4. *Inference and Practice*

As ought to be expected from his discussion of the regress-of-rules argument, it is important to Sellars that the inferential conception of concepts connects the grasp or understanding of concepts (the uptake of conceptual content) with a certain kind of practical activity. Inferring is a kind of doing. Acknowledgment of inferential proprieties need not be explicit in the endorsement of rules or principles of inference but may remain implicit in the capacity to take or treat inferential transitions as correct or incorrect in practice. Inferential relations among concepts are implicit in the practice of giving and asking for reasons. The norms that govern these justificatory practices can be understood to confer inferentially articulated contents on the states, attitudes, and performances subject to them: for something to have such content just is for such norms to determine how it is correctly used or manipulated. The status of inference as something that can be *done* accordingly holds out the promise of securing an appropriate relation between pragmatics, the study of the practices, and semantics, the study of the corresponding contents.

Furthermore, because the activity through which the norms get their grip on conceptual contents is construed as *inference*, it is specifically *propositional* contents that in the first instance count as conceptually articulated. Inferential relations hold, in the paradigm case, between contents that are expressed explicitly by declarative sentences. The premises of inferences, and in the central cases their conclusions as well, must be understood to have a

propositional form. Insofar as an independent theoretical grip is possible on the notion of inference, propositional contents can be picked out by appeal to this property. (This is the strategy pursued in the next chapter.) So on an inferential rendering of the conceptual, the sort of doing that inferring yields in a natural way the priority of propositional conceptual contents.

The pragmatic turn aside, this view too is due to Kant. As Sellars puts it: "Kant was on the right track when he insisted that just as concepts are essentially (and not accidentally) items which can occur in judgments, so judgments (and therefore, indirectly concepts) are essentially (and not accidentally) items which can occur in reasonings or arguments."<sup>37</sup> The subtlety and sophistication of Kant's concept of representation is due in large part to the way in which it is integrated into his account of the inferential relations among judgments. It remained for Hegel, however, to complete the inversion of the traditional order of semantic explanation by beginning with a concept of experience as inferential activity and discussing the making of judgments and the development of concepts entirely in terms of the roles they play in that inferential activity. Although something like this point had been implicit in Kant's notion of reason as systematicity, it was the young Hegel who first appreciated the line of reasoning, made familiar to us by Quine in "Two Dogmas"—namely, that if the content of a claim must at least determine what follows from it (what else it commits one to), then since what a claim commits one to depends on what collateral commitments are available to serve as additional premises ("auxiliary hypotheses"), the significance of undertaking any particular commitment cannot be determined without appeal to the contents of all those collateral commitments.

Hegel's two central semantic concepts in the *Phenomenology* are both inferential notions. "Mediation," his term for inferential articulation, is derived from the role of the middle term in syllogistic inference. "Determinate negation" is his term for material incompatibility, from which, he takes it, the notion of formal negation is abstracted. The contents of concepts are identified and individuated by the functional roles they play in historically evolving webs constituted by relations of mediation and determinate negation, that is, by their material inferential and incompatibility relations to each other. Hegel's interest in the significance of inference in semantics does not (as with Kant) arise primarily in the investigation of how it might be combined with representationalist insights (although he has something to say about that too). It arises rather in the investigation of how this rationalist insight might be combined with the insights of the Romantic expressivists.

The Romantics are perhaps best known for their rejection, not just of the Enlightenment's representationalism, discussed above, but also for their rejection of the significance it assigns to reason. They sought to displace the general demarcational emphasis on giving and asking for reasons or inquiring after truth, not just the specific version that sought to understand these matters in representational terms. The Romantic recoil from understanding us as representers overshoot that mark and came to rest in an esteem for

feeling and inarticulate empathy and enthusiasm. Hegel saw in inferential notions of content a way to join the Romantics in rejecting representationalism, while parting company with them in their hostility to reason. The result is a synthesis of Enlightenment inferentialism and Romantic expressivism.<sup>38</sup>

### 5. *Inferentialism and Representationalism*

Kant, however, did not originate the inferentialist line of semantic thought that Sellars appropriates from him, and that Hegel develops. In a discussion of his break with traditional empiricism, prompted by the issue of the sort of content that ought to be associated with logical, causal, and deontological modalities, Sellars puts the idea that stands at the center of his systematic thought in the form in which it originally occurred to him in the 1930s: "What was needed was a functional theory of concepts which would make their role in reasoning, rather than supposed origin in experience, their primary feature."<sup>39</sup> Put this way, the idea forms one of the mainstays of classical rationalism, even in the absence of Kant's insight about the privileged role that must be assigned to judgments on such an inferential-functional approach. Pre-Kantian empiricists and rationalists alike were notoriously disposed to run together causal and conceptual issues, largely through insufficient appreciation of the normative character of the "order and connection of ideas" that matters for the latter. But there is another, perhaps less appreciated, contrast at work here, besides that of the causal and the conceptual. Enlightenment epistemology was always the home for two somewhat uneasily coexisting conceptions of the conceptual. The fundamental concept of the dominant and characteristic understanding of cognitive contentfulness in the period initiated by Descartes is of course *representation*. However there is a minority semantic tradition that takes *inference* rather than representation as its master concept.

Rationalists such as Spinoza and Leibniz accept the central role of the concept of representation in explaining human cognitive activity, but they are not prepared to accept Descartes's strategy of treating the possession of representational content as an unexplained explainer. Each of them develops instead an account of what it is for one thing to represent another, in terms of the *inferential* significance of the representing. They are explicitly concerned (as Descartes is not) to be able to explain what it is for something to be understood, treated, or employed in practice *as* a representing *by* the subject—what it is for it to be a representing *to* or *for* that subject. Their idea is that the way in which representings point beyond themselves to something represented is to be understood in terms of *inferential* relations among representings. States and acts acquire content by being caught up in inferences, as premises and conclusions.<sup>40</sup> Thus a big divide within Enlightenment epistemology concerns the relative explanatory priority accorded to the concepts of representation and inference.

The British empiricists are more puzzled than Descartes about representational purport, the property of seeming to be about something. But they are clear in seeking to derive inferential relations from the contents of representings, rather than the other way around. In this regard they belong to the still-dominant tradition that reads inferential correctnesses off from representational correctnesses, which are assumed to be antecedently intelligible. The post-Cartesian rationalists, the claim is, give rise to a tradition based on a complementary, semantically reductive order of explanation. These *inferentialists* seek to define representational properties in terms of inferential ones, which must accordingly be capable of being understood antecedently. They start with a notion of content as a matter of what is a reason for what and understand truth and representation as features of ideas that are not only manifested in, but conferred by their role in reasoning. This is the tradition that Sellars inherits and builds on by developing a notion of conceptual content that starts with inferential roles.

#### IV. MATERIAL INFERENCE, CONCEPTUAL CONTENT, AND EXPRESSION

##### 1. Frege on Conceptual Content

The rationalists' inferential understanding of conceptual content, which Kant inherits and which remains one of the strands from which his systematic semantic tapestry is woven, provides the starting point as well for Frege's semantic investigations. Frege may seem an unlikely heir to this inferentialist tradition. After all, he is usually thought of as the father of the contemporary way of working out the representationalist order of explanation. Its strategy is to start with an independent notion of relations of reference or denotation obtaining between mental or linguistic items and objects and sets of objects in the largely nonmental, nonlinguistic environment. Then it determines from these in the familiar fashion: first truth conditions for the sentential representings built out of the subsentential ones, and then, from these, a notion of goodness of inference understood in terms of set-theoretic inclusions among the associated sets of truth conditions. But insofar as it is appropriate to read this twentieth-century story back into Frege at all (a dangerous and potentially misleading enterprise), it would be possible only beginning with the Frege of the 1890s. He starts his semantic investigations, not with the idea of *reference*, but with that of *inference*. His seminal first work, the *Begriffsschrift* of 1870, takes as its task the explicit expression of *inferential* roles: "In my formalized language [*Begriffsschrift*] . . . only that part of judgments which affects the possible inferences is taken into consideration. Whatever is needed for a correct [*richtig*] inference is fully expressed; what is not needed is . . . not."<sup>41</sup>

These inferential roles form the basis of his notion of *content*. It is because the sorts of contents that are associated with expressions are to be defined

in the first place in terms of inference that Frege must insist on the distinction between the sorts of contents that can, and those that cannot, serve as premises and conclusions of inference, and so play the basic sort of inferential roles. "We distinguish contents that *are*, and contents that *are not, possible contents of judgment*."<sup>42</sup> Frege's Kantian insistence on the priority of the propositional, of judgeable contents, is an aspect of his pursuit of the rationalists' inferentialist order of semantic explanation. He embraces Kant's insight that the notion of content must be made intelligible first for judgments, which alone can figure as premises and conclusions of inference, and only then extended to the contents expressed by fragments of declarative sentences. Recall the passage (already quoted in Section II of this chapter) in which he contrasts his procedure with that pursued by others in the tradition: "In Aristotle, as in Boole, the logically primitive activity is the formation of concepts by abstraction, and judgment and inference enter in through an immediate or indirect comparison of concepts via their extensions . . . I start out from judgments and their contents, and not from concepts . . . I only allow the formation of concepts to proceed from judgments . . . Instead of putting a judgment together out of an individual as subject and an already previously formed concept as predicate, we do the opposite and arrive at a concept by splitting up the content of a possible judgment."<sup>43</sup> It is for this reason that the fundamental definition introducing the notion of "conceptual content" (*begriffliche Inhalt*) (for which, as its name implies, the *Begriffsschrift* is supposed to supply a means of explicit expression) applies only to the contents of possible judgments. It will have to be extended later, by Frege's substitutional methodology, to allow the assignment of *indirectly* inferential roles to *subsential* expressions, according to the contribution their occurrence makes to the directly inferential role (as premise or conclusion) of judgment-expressing sentences in which they occur. The substitutional strategy that Frege devised for quarrying subsentially expressed contents from sententially expressed ones is of the first importance for carrying out the inferentialist semantic explanatory program. Much is made of it in subsequent chapters of this work. Before Frege, one could only hope that there was some way of bridging this gap.

That the target notion of content is specifically *conceptual* content is accordingly not to be understood in terms of some antecedent notion of concepts. Rather, the conceptual is explicitly construed in *inferential* terms: "There are two ways in which the content of two judgments may differ; it may, or it may not, be the case that all inferences that can be drawn from the first judgment when combined with certain other ones can always also be drawn from the second when combined with the same other judgments. The two propositions 'the Greeks defeated the Persians at Plataea' and 'the Persians were defeated by the Greeks at Plataea' differ in the former way; even if a slight difference of sense is discernible, the agreement in sense is preponderant. Now I call that part of the content that is the same in both the

conceptual content. Only this has significance for our symbolic language [*Begriffsschrift*]."<sup>44</sup> Two claims have the same conceptual content if and only if they have the same inferential role: a good inference is never turned into a bad one by substituting one for the other. The fundamental semantic assignment of conceptual content to judgments is derived from the ultimately pragmatic notion of *correctness* of *inference*. This derivation is the first application of the substitutional methodology: semantically assimilating expressions accordingly as substitution of one for another preserves some semantically relevant property. In this case (prior to the others in the order of explanation), the semantically relevant invariant is propriety of inference.

This way of specifying the explanatory target to which semantic theories, including referential ones, are directed is picked up by Frege's student Carnap, who in *The Logical Syntax of Language* defines the content of a sentence as the class of nonvalid sentences that are its consequences (that is, can be inferred from it). Sellars in turn picks up the idea from him, as his references to this definition indicate.<sup>45</sup> As will emerge, an important feature of Carnap's definition is the appeal to *nonvalid* consequences. In this way what pertains to the *content* of a claim is distinguished from what pertains to its *form*.

This distinction is operative in the *Begriffsschrift* as well. Yet when Frege wants to be clear about what is expressed by even the purely formal assertions appearing in proofs about the expressive capacity of the *Begriffsschrift* itself, he does so by specifying their inferential role, restricting himself in this case to inferences whose propriety is underwritten by their form alone. So each assertion is introduced by displaying a proof of it from already-established assertions, thereby exhibiting the premises from which it follows as conclusion. Showing what a claim follows from is not sufficient to specify its inferential role, however. It matters as well what follows from it. Indeed, Frege often complains (for instance in the *Grundlagen*) about systems that introduce definitions that are never then employed in subsequent demonstrations. These provide a case where looking at inferential consequences is particularly important; since definitions do not have inferential antecedents, if their inferential consequents are not specified, their content is left entirely indeterminate. In order to complete the specification of the inferential roles of the assertions of the system he presents, Frege appends to the *Begriffsschrift* a list indicating for each assertion all of the subsequent assertions in whose proof it is used as premise. That is, he specifies for each assertion what follows from it (together with other assertions, of course) as well as what it follows from. In this way he makes explicit the inferential roles, and so the conceptual contents, conferred on the judgments he puts forward by the purely formal reasoning involving them that is displayed in his book.

In contrast to his original procedure, the tradition Frege initiated in the 1890s makes truth, rather than inference, primary in the order of semantic explanation. Dummett says of this shift: "In this respect (and in this respect



alone) Frege's new approach to logic was retrograde. He characterized logic by saying that, while all sciences have truth as their goal, in logic truth is not merely the goal, but the object of study. The traditional answer to the question what is the subject-matter of logic is, however, that it is, not truth, but inference, or, more properly, the relation of logical consequence. This was the received opinion all through the doldrums of logic, until the subject was revitalized by Frege; and it is, surely, the correct view."<sup>46</sup> And again: "It remains that the representation of logic as concerned with a characteristic of sentences, truth, rather than of transitions from sentences to sentences, had highly deleterious effects both in logic and in philosophy. In philosophy it led to a concentration on logical truth and its generalization, analytic truth, as the problematic notions, rather than on the notion of a statement's being a deductive consequence of other statements, and hence to solutions involving a distinction between two supposedly utterly different kinds of truth, analytic truth and contingent truth, which would have appeared preposterous and irrelevant if the central problem had from the start been taken to be that of the character of the relation of deductive consequence."<sup>47</sup> The important thing to realize is that the Frege of the *Begriffsschrift* has not yet made this false step. Of course, adopting a semantic order of explanation that begins with proprieties of inference requires both an account of those proprieties (that is, an account of its raw materials) and an account of how talk about truth is eventually to be construed in these terms (that is, an account of its consequences). This is the strategy pursued in this work. The first of these challenges is responded to in Chapters 3 and 4, and the second in Chapter 5.

There are two further points to keep in mind regarding this passage of Dummett's. First, shifting from concern with *inference* to concern with *truth* is one move; understanding *truth* in terms of prior primitive *reference* relations involving objects and properties is another. Since the mature Frege treats truth as indefinable and primitive, the extraction of a representationalist commitment even from the texts of the 1890s requires further showing (compare Davidson's truth-without-reference view in our own day). Second, understanding the topic of logic in terms of inference is not the same as seeing it in terms of *logical* inference, or of "deductive consequence," as Dummett puts it (see the discussion of "formalism" about inference, below). The view propounded and attributed to Frege below is a different one—and from the contemporary vantage point it is a more surprising one than the one that Dummett endorses here.

## 2. *Material Proprieties of Inference and the Dogma of Formalism*

The kind of inference whose correctnesses essentially involve the conceptual contents of its premises and conclusions may be called, following Sellars, "*material* inference."<sup>48</sup> As examples, consider the inference from

"Pittsburgh is to the West of Philadelphia" to "Philadelphia is to the East of Pittsburgh," the inference from "Today is Wednesday" to "Tomorrow will be Thursday," and that from "Lightning is seen now" to "Thunder will be heard soon." It is the contents of the concepts *West* and *East* that make the first a good inference, the contents of the concepts *Wednesday*, *Thursday*, *today*, and *tomorrow* that make the second inference correct, and the contents of the concepts *lightning* and *thunder*, as well as the temporal concepts, that underwrite the third. Endorsing these inferences is part of grasping or mastering those concepts, quite apart from any specifically logical competence. From the point of view of a familiar sort of semantics (different from that to be explored here), one could say that the set of possible worlds in which the premises of these inferences are true is a subset of the set of possible worlds in which their conclusions are true. Since neither the premises nor the conclusions of such inferences employ logical concepts, it seems appropriate to distinguish them from inferences whose correctness depends only on logical form.

Often, however, *inferential* articulation is identified with *logical* articulation. Material inferences are then treated as a derivative category. The idea is that being *rational*—mastering proprieties of inference and so being subject to the force of the better reason—can be understood as a purely *logical* capacity. In part this tendency is encouraged by merely verbally sloppy formulations of the crucial difference between the inferential force of reasons and the physically efficacious force of causes: formulations that render it as the difference between 'logical' and 'natural' compulsion. Mistakes ensue, however, if the concept *logical* is employed with these circumstances of application conjoined with consequences of application that restrict the notion of the logical force of reasons to *formally* valid inferences. The substantial commitment that is fundamental to this sort of approach is what Sellars calls "the received dogma . . . that the inference which finds its expression in 'It is raining, therefore the streets will be wet' is an enthymeme."<sup>49</sup>

According to this line of thought, wherever an inference is endorsed, it is because of belief in a conditional. Then the instanced inference is understood as implicitly involving the conditional "If it is raining, then the streets will be wet." With that "suppressed" premise supplied, the inference is an instance of the formally valid scheme of conditional detachment. The "dogma" expresses a commitment to an order of explanation that treats all inferences as good or bad solely in virtue of their form, with the contents of the claims they involve mattering only for the truth of the (implicit) premises. According to this way of setting things out, there is no such thing as material inference. This view—which understands "good inference" to mean "formally valid inference," postulating implicit premises as needed—might be called a *formalist* approach to inference. It trades primitive goodnesses of inference for the truth of conditionals. Doing so is taking a retrograde step that corresponds to the one Dummett complains about. The grasp of logic

that is attributed must be an *implicit* grasp, since it need be manifested only in distinguishing material inferences as good and bad, not in any further capacity to manipulate logical vocabulary or endorse tautologies involving them. But what then is the explanatory payoff from attributing such an implicit logical ability rather than just the capacity to assess proprieties of material inference?

It is worth considering an example of how formalist presuppositions can be embodied misleadingly in vocabulary. Here is Dennett in "Intentional Systems":

Earlier I alleged that even creatures from another planet ["in virtue of their rationality"] would share with us our beliefs in logical truths; light can be shed on this claim by asking whether mice and other animals, in virtue of being intentional systems, also believe the truths of logic. There is something bizarre in the picture of a dog or mouse cogitating a list of tautologies, but we can avoid that picture. The assumption that something is an intentional system is the assumption that it is rational; that is, one gets nowhere with the assumption that entity  $x$  has beliefs  $p, q, r \dots$  unless one also supposes that  $x$  believes what follows from  $p, q, r \dots$ ; otherwise there is no way of ruling out the prediction that  $x$  will, in the face of beliefs  $p, q, r \dots$  do something utterly stupid, and, if we cannot rule out that prediction, we will have acquired no predictive power at all. So whether or not the animal is said to believe the truths of logic, it must be supposed to follow the rules of logic.<sup>50</sup>

Dennett understands intentionality in terms of rationality (as the view being developed here does), and understands rationality in terms of the discrimination in practice of good inferences ("what follows") from bad ones (as the view being developed here does). But there is a slide here from "follows" to "logically follows." No justification is offered for the move, first, from discriminating good from bad inferences to the need for any specifically logical capacity or, second, for the move from logical capacity to belief in logical truths. Perhaps appropriate (even logically valid) inferences can be endorsed without commitment to the corresponding (logical) conditional truths.

On the first point: perhaps there are good nonlogical inferences, and rationality consists in the way discriminating them matters to one's deliberations and assessments. Why should "following the rules of logic" be either necessary or sufficient for this discrimination? In any case, it was argued in Chapter 1 that one ought to distinguish both exhibiting a regularity and acknowledging a norm implicitly in one's practice (two construals of discriminating good from bad inferences) from following a rule. On the second point, Dummett was cited above as pointing out that defining logical consequence in terms of logical truth is neither a trivial nor a harmless move.

In fact Dennett (and in this regard he is typical) thinks of this way of putting things as a harmless *façon de parler*, warranted by a general inter-

changeability of talk of endorsing inferences and talk of believing conditionals. The "belief in logical truths," or even, less committally, endorsement of logically good inferences, that he has in mind is implicit in practical discriminations. The passage continues: "Surely our mouse follows or believes in *modus ponens*, for we ascribed to it the beliefs: (a) there is a cat to the left, and (b) if there is a cat to the left, I had better not go left, and our prediction relied on the mouse's ability to get to the conclusion." What was actually attributed to the mouse is a belief with content (a) and a desire to avoid the cat. Citing its intelligent behavior licenses the attribution of a practical inference. It does not, by itself, tell for or against expressing that inference as a material inference or as detachment from an endorsed conditional. Why should all inferences be assimilated to detachments, or other formal logical rules of inference? Dennett's justification is that "in general there is a trade-off between rules and truths; we can suppose *x* to have an inference rule taking *A* to *B* or we can give *x* the belief in the 'theorem': if *A* then *B*. As far as our predictions are concerned, we are free to ascribe to the mouse either a few inference rules and belief in many logical propositions, or many inference rules and few if any logical beliefs."<sup>51</sup>

The conditional beliefs that can be traded off for endorsements of inferences should not be called "logical" beliefs simply because they concern inferences. Though they involve logical concepts, namely the conditional, they are not in general logically true. Indeed, Dennett continues: "We can even take a patently nonlogical belief like (b) and recast it as an inference rule taking (a) to the desired conclusion." To do so would be to establish or endorse a *material* correctness of inference, what Dennett calls "a set of nonlogical inference rules." Once the possibility of this sort of inference is acknowledged, inferential formalism surrenders a priori privileges and must contend with inferential materialism for privileges of explanatory priority. According to the famous argument of Lewis Carroll in "What the Tortoise Said to Achilles," as Dennett acknowledges, some inferential commitments ("rules of inference") must be attributed if any consequences are to be licensed by the attribution of beliefs, even conditional beliefs. So there must be "rules" as well as "truths." However, once the purely formal-logical inferences are allowed (paradigmatically detachment inferences licensed by conditionals), accounts of rational performance can take the form either of attributions of endorsements of material inferences or of conditional propositions, as might be theoretically convenient for other reasons. Either decision ought to be justified.

What considerations ought to persuade a theorist to accord explanatory priority to the attribution of material inferential commitments or to the attribution of conditional propositional commitments, and so to treat material or formal inference as fundamental? Dennett's answer is: "If we found an imperfectly rational creature whose allegiance to *modus ponens*, say, varied with the subject matter, we could characterize that by excluding *modus*

ponens as a rule and ascribing in its stead a set of nonlogical inference rules covering the modus ponens step for each subject matter where the rule was followed."<sup>52</sup> This is a formalist position, in that all inferences are assimilated to detachments and are understood as involving, at least implicitly, endorsements of conditionals whose logical content explicitly relates premises and conclusions. The only concession to material inferences arises in the possibility of licensing detachment in a retail, content-respecting fashion, rather than wholesale, in a purely formal logical way. But why should this model be employed? Why should all goodness of inference be seen as logical goodness, even at the cost of postulating "implicit" premises involving logical concepts?

What is at issue is two different ways of understanding the relation between something implicit and an explicit expression of it. It is possible to agree with the formalist in understanding conditionals as inference licenses, which make explicit in the content of a claim what is implicit in the endorsement of an inference, without going on to construe all inferences as involving the use of conditionals. The question is how one ought to construe the relation between what is explicit in the form of a rule or principle (in this case a conditional claim) and what is implicit in proprieties of practice (in this case in the endorsement of an inference). The formalist line of thought begins with explicit propositional licenses that license inferences in virtue of their logical form. Material inferences (say from rain to wet streets or vice versa) are understood privatively: as enthymemes resulting from the suppression or hiding of one of the premises required for a proper warrant. Opposed to this might be a pragmatist line of thought, beginning with material inferences—that is, nonlogical, content-based reasoning. It would then be necessary to explain how logical vocabulary such as the conditional is to be understood as permitting the expression of those implicit inferential commitments in an explicit fashion—that is, as judgeable, claimable, believable contents, as the contents of potential propositional commitments.

There are general reasons to prefer an order of explanation that begins with what is implicit in practice (what people do) and proceeds to an account of what they explicitly believe or say, over one taking the opposite tack. Only in this way can one hope to understand believing or saying in terms of more primitive capacities (knowing-that in terms of knowing how). That asymmetry manifests itself in this case in the question of how one understands logical concepts or the use of logical vocabulary. On the formalist line, anything that has any inferential capacities at all is credited with mastery of a battery of logical concepts and the corresponding inference rules without which they would be without content. These can be thought of as introduction and elimination rules, of which detachment is a cardinal example. Logical concepts are quite different from others in being presupposed by all contentful concepts and inferences. It is a short step from treating mastery of these concepts as implicit in inferential abilities to treating it as an innate

presupposition of them. This sort of thing gave the classical rationalists a bad name. Kant rescued them by insisting that it is the formality of logical (and, more controversially, transcendental) concepts that entitles them to a special status that would indeed be absurd for ordinary contentful concepts. Assessing the Kantian formalist move requires looking more closely at what is being said when an inference is described as being valid in virtue of its logical form.

### 3. *Conceptual Content and Material Inference*

Before looking at how logical concepts might function to make explicit conceptual contents that are implicit in practical proprieties of inference, however, it is worth looking more closely at the relation between inference and content. The picture being developed is one according to which materially good inferences correspond to the conceptual content of nonlogical expressions, while inferences valid in virtue of their logical form alone correspond to the conceptual content of purely logical expressions. This can be approached by considering, to begin with, the notion of *material* inferences: inferences whose propriety essentially involves the nonlogical conceptual content of the premises and conclusions. The approach Sellars endorses is best understood by reference to the full list of alternatives he considers:

We have been led to distinguish the following six conceptions of the status of material rules of inference:

(1) Material rules are as essential to meaning (and hence to language and thought) as formal rules, contributing to the architectural detail of its structure within the flying buttresses of logical form.

(2) While not essential to meaning, material rules of inference have an original authority not derived from formal rules, and play an indispensable role in our thinking on matters of fact.

(3) Same as (2) save that the acknowledgment of material rules of inference is held to be a dispensable feature of thought, at best a matter of convenience.

(4) Material rules of inference have a purely derivative authority, though they are genuinely rules of inference.

(5) The sentences which raise these puzzles about material rules of inference are merely abridged formulations of logically valid inferences. (Clearly the distinction between an inference and the formulation of an inference would have to be explored.)

(6) Trains of thought which are said to be governed by "material rules of inference" are actually not inferences at all, but rather activated associations which mimic inference, concealing their intellectual nudity with stolen "therefores."<sup>53</sup>

His own position is that an expression has conceptual content conferred on it by being caught up in, playing a certain role in, material inferences: "It is

the first (or 'rationalistic') alternative to which we are committed. According to it, material transformation rules determine the descriptive meaning of the expressions of a language within the framework provided by its logical transformation rules . . . In traditional language, the 'content' of concepts as well as their logical 'form' is determined by the rules of the Understanding."<sup>54</sup>

Sellars, in arguing that material inferences are essential to the meaning (content) of nonlogical locutions, cites a phenomenon that is as important to the expressivist picture of logical concepts as it is to the materialist conception of inference presupposed by inferentialist approaches to conceptual content. Sellars's argument that material inferences are essential to the meaning (content) of nonlogical locutions depends on a central conceptual phenomenon. He argues for the theoretical indispensability of a conception of material inferences in terms of the practical indispensability of what is made explicit by a certain familiar kind of vocabulary. His argument is attributed to an interlocutor who maintains that:

such subjunctive conditionals as "If I had released this piece of chalk, it would have fallen," and "If there were to be a flash of lightning, there would be thunder" . . . [must be interpreted] as expressions of material rules of inference . . . He therefore claims to have shown beyond reasonable doubt not only that there are such things as material rules of inference, but, which is far more important, that they are essential to any conceptual frame which permits the formulation of such subjunctive conditionals as do not give expression to logical principles of inference. Since we are all conscious of the key role played in the sciences, both formal and empirical, in detective work and in the ordinary course of living by subjunctive conditionals, this claim, if substantiated, would indeed give a distinguished status to material rules of inference.<sup>55</sup>

He concludes: "Now, unless some other way can be found of interpreting such subjunctive conditionals in terms of logical principles of inference, we have established not only that they are the expression of material rules of inference, but that the authority of these rules is not derivative from formal rules. In other words, we have shown that material rules of inference are essential to the language we speak, for we make constant use of subjunctive conditionals."<sup>56</sup>

The point is not the indispensability of the vocabulary of conditionals that permit detachment inferences even with counterfactual premises. It is the indispensability of what those conditionals express: the implicit proprieties of material inference that they help make explicit. "Even though material subjunctive conditionals may be dispensable, permitting the language to be extensional, it may nevertheless be the case that the function performed in natural languages by material subjunctive conditionals is indispensable."<sup>57</sup>

The material inferences codified in subjunctive conditionals are inferential involvements that are essential to the contents of the concepts used in science and everyday life. These are not logically valid inferences. But logical

vocabulary, subjunctive conditionals, can be used to express these material inferential relations. Without such vocabulary, the inferences can still be endorsed. With it, those content-generating inferential endorsements can be made explicit as the content of a claim or propositional endorsement.

#### 4. *From Material to Formal Proprieties of Inference*

Should inferentialist explanations begin with inferences pertaining to propositional *form*, or those pertaining to propositional *content*? One important consideration is that the notion of formally valid inferences is definable in a natural way from that of materially correct ones, while there is no converse route. For given a subset of vocabulary that is privileged or distinguished somehow, an inference can be treated as good in virtue of its form, with respect to that vocabulary, just in case it is a materially good inference and cannot be turned into a materially bad one by substituting nonprivileged for privileged vocabulary, in its premises and conclusions. This is another application of the substitutional methodology Frege employs in individuating the conceptual contents of judgments, and again in discerning indirectly conceptually contentful components within them. All it requires is a partition of vocabulary into two kinds: those that are to be held fixed and those that are to be regarded as replaceable. Call the kind of vocabulary that is to be held fixed the *K-vocabulary*. The general structure of formality definitions is then that the set of *K-valid* inferences (those that will be understood as good in virtue of their *K-form* alone) comprises those that meet the two conditions of being inferences that (1) are good inferences and (2) cannot be turned into bad inferences by substituting non-*K* for *K* vocabulary.

Clearly, what inferences are treated as valid in virtue of their form by such a procedure depends on how the vocabulary is divided into the two kinds. In the limit, if *all* the vocabulary were treated as irreplaceable, no substitutions of non-*K* for *K* vocabulary would be possible, and a fortiori none could turn a correct inference into one that is not correct. So *all* materially good inferences would count as good in virtue of their *K-form*, in the case where *K* comprises the whole vocabulary over which the field of inferences is defined. At the opposite end of the spectrum, if *no* vocabulary is treated as irreplaceable, then if there were any bad inferences at all, none of the good inferences would count as good in virtue of their *K-form*. For all could be turned into bad inferences by some substitution or other.

If the *K-vocabulary* (that which is not substituted for) is *logical* vocabulary, then the good inferences whose correctness is invariant under substitution of non-*K* for *K* vocabulary (nonlogical for logical vocabulary) are the *logically* valid inferences—namely those that are good in virtue of their *logical* form. (Quine recommends this Fregean substitutional way of thinking about logical form, although he appeals to *truth* rather than propriety of



*inference* as the semantically relevant invariant whose preservation is at issue.) But this substitutional conception of what it is for an inference to be good in virtue of its form is not essentially restricted to a notion of *logical* form. If one picks out specifically zoological vocabulary or moral vocabulary or theological vocabulary to play the role of the distinguished *K*-vocabulary, the substitutional mechanism will take as its input a practical classification of inferences into good or bad, correct or incorrect, and yield as its output a distinguished set of inferences that are not just good, but are good in virtue of their zoological, moral, or theological form. The mechanism is perfectly general.

It follows that on this way of thinking about things, logical vocabulary cannot be picked out by appeal to its formality or by its involvement in formal proprieties of inference. If it is specifically logical form that is of interest, then one must antecedently be able to distinguish some vocabulary as peculiarly logical. That done, the Fregean semantic strategy of looking for inferential features that are invariant under substitution yields a notion of logically valid inferences. So the formal goodness of inferences derives from and is explained in terms of the material goodness of inferences, and so ought not to be appealed to in explaining it. And logical vocabulary must be picked out in some way that does not appeal to inferences that are formally valid or good in virtue of their form. Frege's way of specifying the characteristic linguistic role in virtue of which vocabulary qualifies as logical is discussed below.

### 5. Sellars on Expressive Rationality

So far two related claims have been introduced: that conceptual contents are inferential roles, and that the inferences that matter for such contents in general must be conceived to include those that are in some sense materially correct, not just those that are formally valid.<sup>58</sup> It will be argued in a moment that a commitment to the second of these, no less than the first, is to be found already in Frege's early writings, though not in the developed form to which Sellars brings it. But in both thinkers these ideas are combined with a third, which makes this line of thought especially attractive. In one of his early papers, Sellars introduces the idea this way: "In dealing with such situations [attempts to justify acceptance of a law by means of an argument from instances], philosophers usually speak of inductive arguments, of establishing laws by induction from instances . . . I am highly dubious of this conception. I should be inclined to say that the use Jones will make of instances is rather in the nature of Socratic method. For Socratic method serves the purpose of making explicit the rules we have adopted for thought and action, and I shall be interpreting our judgments to the effect that *A* causally necessitates *B* as the expression of a rule governing our use of the terms '*A*' and '*B*'."<sup>59</sup> Sellars understands such modal statements as inference

licenses, which formulate as the content of a claim the appropriateness of inferential transitions. More than this, he understands the function of such statements to be making explicit, in the form of assertible rules, commitments that had hitherto remained implicit in inferential practices. Socratic method is a way of bringing our practices under rational control, by expressing them explicitly in a form in which they can be confronted with objections and alternatives, a form in which they can be exhibited as the conclusions of inferences seeking to justify them on the basis of premises advanced as reasons, and as premises in further inferences exploring the consequences of accepting them.

In the passage just cited, Sellars tells us that the enterprise within which we ought to understand the characteristic function of inductive inference is a form of rationality that centers on the notion of expression: making explicit, in a form that can be thought or said, what is implicit in what is done. This is a dark and pregnant claim, but it epitomizes a radical and distinctive insight. What follows is intended to shed some light on it and its role in an inferentialist vision of things. The general idea is that the paradigmatically rational process that Sellars invokes under the heading of "Socratic method" depends upon the possibility of making *implicit* commitments *explicit* in the form of claims. Expressing them in this sense is bringing them into the game of giving and asking for reasons as playing the special sort of role in virtue of which something has a conceptual content at all—namely an inferential role, as premise and conclusion of inferences.

This is distinct from (but obviously related to) the sort of rationality that then consists in making the appropriate inferential moves. Even totalitarian versions of the latter—for instance those that would assimilate all goodness of inference to logical validity, or to instrumental prudence (that is, efficiency at getting what one wants)—depend upon the possibility of expressing considerations in a form in which they can be given as reasons, and reasons demanded for them. All the more does Socratic reflection on our practices, particularly on those material-inferential practices that determine the conceptual contents of thoughts and beliefs, depend on the possibility of their explicit expression. Here is another early (perhaps equally dark) statement of this important Sellarsian theme:

Now, among the linguistic activities which can be discriminated are the 'explicative' or 'analytic' which, to use Ayer's phrase 'elucidate the proper use' of linguistic expressions. Furthermore the anthropologist . . . can distinguish within language activity between that which "deals directly with the environment" and that which attempts to mirror, within language itself, the relation of language to the world. In connection with this Fichtean self-diremption, the language user makes use of such words as 'means', 'true', 'verified' and so on. This is linguistic activity as semantic and pragmatic metalanguage. But the language

activity of human organisms can achieve an even greater degree of internal complexity, such as comes out most clearly in the 'explicative' metalinguistic activity of the logician and epistemologist, but is also to be found, highly confused, in more practical beings.<sup>60</sup>

### 6 . *The Expressive Project of the Begriffsschrift*

To begin to explicate this notion of explication, it is helpful to return to the consideration of the young Frege's inferentialist program. Frege's *Begriffsschrift* is remarkable not only for the inferential idiom in which it specifies its topic but equally for how it conceives its relation to that topic. The task of the work is officially an *expressive* one—not to *prove* something, but to *say* something. Frege's logical notation is designed for *expressing* conceptual contents, making *explicit* the inferential involvements that are *implicit* in anything that possesses such content. As the passage quoted above puts it: "Whatever is needed for a correct inference is fully expressed." Talking about this project, Frege says: "Right from the start I had in mind the expression of a content . . . But the content is to be rendered more exactly than is done by verbal language . . . Speech often only indicates by inessential marks or by imagery what a concept-script should spell out in full."<sup>61</sup> The concept-script is a formal language for the explicit codification of conceptual contents. In the preface to the *Begriffsschrift*, Frege laments that even in science, concepts are formed haphazardly, so that the ones employing them are scarcely aware of what they mean, of what their content really is. When the correctness of particular inferences is at issue, this sort of unclarity may preclude rational settlement of the issue. What is needed, he thinks, is a notation within which the rough-and-ready conceptual contents of the sciences, beginning with mathematics, can be reformulated so as to wear their contents on their sleeves. His explanatory target avowedly concerns a sort of inference, not a sort of truth, and the sort of inference involved must be content-conferring material inferences, whose propriety is determined before logical vocabulary comes on the scene, not the derivative formal ones whose propriety is underwritten by the use of that vocabulary.

Frege explicitly contrasts his approach with that of those, such as Boole, who conceive their formal language only in terms of formal inference, and so express no material contents: "The reason for this inability to form concepts in a scientific manner lies in the lack of one of the two components of which every highly developed language must consist. That is, we may distinguish the formal part . . . from the material part proper. The signs of arithmetic correspond to the latter . . . In contrast, Boole's symbolic logic only represents the formal part of the language."<sup>62</sup> Frege's own project is to express the contents that make up the material part of the language, not just the "formal cement that can bind these stones together": "My concept-script has a more far-reaching aim than Boolean logic, in that it strives to make it

possible to present a content when combined with arithmetical and geometrical signs . . . It is in a position to represent the formation of the concepts actually needed in science."<sup>63</sup> It is the wider domain to which his expressive ambition extends that Frege sees as characteristic of his approach.

Since contents are determined by inferences, expressing inferences explicitly will permit the expression of any sort of content at all: "It seems to me to be easier still to extend the domain of this formula language to include geometry. We would only have to add a few signs for the intuitive relations that occur there . . . The transition to the pure theory of motion and then to mechanics and physics could follow at this point."<sup>64</sup> Indeed, he goes on to suggest that for this reason, "by laying bare the misconceptions that through the use of language often almost unavoidably arise concerning the relations between concepts and by freeing thought from that with which only the means of expression of ordinary language, constituted as they are, saddle it . . . my ideography [*Begriffsschrift*], further developed for these purposes, can become a useful tool for the philosopher."<sup>65</sup>

### 7. Frege's Expressive Conception of Logic

Frege's early understanding of logic offers some specific content to the notion of explicitly expressing what is implicit in a conceptual content. That is what is required to fill in a notion of expressive or elucidating rationality that might be laid alongside (and perhaps even be discovered to be presupposed by) notions of rationality as accurate representation, as logically valid inference, and as instrumental practical reasoning. Before he makes the fateful step from seeing logic as an attempt to codify inferences to seeing it as the search for a special kind of truth (which Dummett bemoans, and to which we owe much of contemporary logic), Frege's aim is to introduce vocabulary that will let one say (explicitly) what otherwise one can only do (implicitly). Consider the conditional, with which the *Begriffsschrift* begins. Frege says of it: "The precisely defined hypothetical relation between contents of possible judgments has a similar significance for the foundation of my concept-script to that which identity of extensions has for Boolean logic."<sup>66</sup> Prior to the introduction of such a locution, one could *do* something, one could treat a judgment as having a certain content (implicitly attribute that content to it), by endorsing various inferences involving it and rejecting others. After conditional locutions have been introduced, one can *say*, as part of the content of a claim, that a certain inference is acceptable. One is able to make explicit material inferential relations between an antecedent or premise and a consequent or conclusion. Since according to the inferentialist view of conceptual contents, it is these implicitly recognized material inferential relations that conceptual contents consist in, the conditional permits such contents to be explicitly expressed. If there is a disagreement about the goodness of an inference, it is possible to say what the dispute

is about and to offer reasons one way or the other. The conditional is the paradigm of a locution that permits one to make inferential commitments explicit as the contents of judgments.

The conditional ("the precisely defined hypothetical relation between contents of possible judgments"), rather than inclusion relations among extensions of concepts, plays the central role in Frege's logic because of two cardinal features of his view that distinguish it from the modern set-theoretic interpretations that develop from Boole's approach. First, he understands the content of nonlogical concepts in terms of their inferential role, rather than in terms of their extensions. Second, he understands the task of logical vocabulary to be expressing explicitly what is implicit in those material conceptual contents. What is implicit in those contents, according to the first or inferentialist commitment, is proprieties of inference. Making what follows from what explicit, as itself a *judgeable* content, one that can itself appear as a premise or conclusion in inference, is exactly the job of the conditional.

Frege's overall project for his *Begriffsschrift* is to use conditionals to make it possible to *say* explicitly what the inferential role of ordinary, nonlogical concepts is. Where, as he thinks is often the case in natural language, the content expressed by words is unclear, the project of expressing them explicitly will show where they need or can use clarification. The project is the rectification of concepts: clarifying them by explicating their contents. It is *saying* what their inferential role is: what follows from the applicability of each concept and what its applicability follows from. Employing the explicating logical locutions of which the conditional is the paradigm is to enable what Frege calls "the scientific formation of concepts." Such concepts will wear their contents on their sleeves; the inferential proprieties in virtue of which they mean what they mean are written down for all to read. The particular sciences can then proceed with their reasonings according to the same standards of rigor in the definition and use of their concepts that nineteenth-century mathematics finally came to aspire to. Although the application of this expressive methodology to the special case of mathematics always was closest to Frege's heart and occupied the greatest part of his energies, right from the beginning he had wider expressive ambitions. So the later writings on the sorts of content to be associated with nonmathematical concepts ("On Sense and Reference" prime among them) ought not to be seen to represent any change of interest or detour from his primary project.

Frege is not as explicit about the role of materially correct inferences as Sellars is, but his commitment to the notion is clear from the relation between two of the views that have been extracted from the *Begriffsschrift*: semantic expressivism about logic and inferentialism about content. Expressivism about logic means that Frege treats logical vocabulary as having a distinctive expressive role—namely making explicit the inferences whose goodness is implicit in the conceptual contents of nonlogical concepts. Infer-

entialism about those conceptual contents is taking them to be identified and individuated by their inferential roles. Together these views require that it be coherent to talk about inference prior to the introduction of specifically logical vocabulary, and so prior to the identification of any inferences as good in virtue of their logical form.<sup>67</sup> In the context of an inferential understanding of conceptual contents, an expressivist approach presupposes a notion of nonlogical inference, the inferences in virtue of which concepts have non-logical content. So the early Frege envisages a field of material inferences that confer conceptual content on sentences caught up in them. Although Frege does not offer an explanation of the concept, in the *Begriffsschrift* his expressive, explicating project commits him to something playing the role Sellars later picks out by the phrase “material inference.”

There is a sense, then, in which the early Frege does see endorsement of conditional judgments as *implicit* in endorsement of the correctness of inferences. It is implicit in exactly the sense that what one is committed to by endorsing an inference as correct, and so by associating a certain conceptual content (that is inferential role) with the premises and conclusion, can be *made* explicit by expressing it in the form of a conditional judgment. The point of introducing logical vocabulary is precisely to make it possible to trade hitherto merely implicit inferential commitments for explicit assertional commitments to conditionals. And the payoff from expressing explicitly (in the form of judgments) the content-constitutive commitments that were implicit in prior inferential practice is the clarification and rectification of those conceptual contents. Formalism about inference—denying the existence of materially good inferences by assimilating all good inferences to logically good inferences, understanding all proprieties of inference as always already underwritten by logical form—turns things on their head. It misses the point of the process of explicitation that Frege puts at the center of the logical enterprise. It is a form of intellectualism, platonism, or *regulism* in the sense defined in the previous chapter. For it sees *rules* or *principles* as already standing behind every propriety of (in this case inferential) practice.

Frege’s primary interest is in the process of explicitation: of expressing what is implicit in a practice, formulating it as an explicit rule or principle. This pragmatist project of explaining how knowing-that is founded on knowing-how, of explaining the codification in (conditional) principles of (inferential) practice, is unintelligible from any theoretical standpoint that acknowledges only the explicit form of propriety. Frege’s fundamental insight into the expressive role of logical vocabulary (above all the conditional) is not incompatible with claiming that commitment to a conditional is implicit in endorsement of an inference, provided one is careful about what is meant by ‘implicit’—provided, that is, that it is understood as making reference to the possibility of engaging in the substantive activity of making it explicit in the form of a claim or principle. But if one goes on to treat all reasoning as explicitly involving *detachment* from conditionals, and there-

fore implicit endorsement of logical truths involving conditionals (including tautologies involving nested or iterated conditionals), then the line has been crossed and Frege's expressive insight has been lost. As should become clearer from the discussion of Section V below, one of the casualties of the inferential formalist's inversion of the significance of the role of conditionals in making explicit in the form of a principle what is implicit in an inferential practice is a proper understanding of the way in which the contents of inferentially articulated concepts evolve and are clarified as they are expressed with the help of logical locutions.

### 8. Expressive Completeness and the Two-Valued Conditional

Various special features of Frege's presentation of his conditional, and of the use he goes on to make of it in the *Begriffsschrift*, tend to obscure the crucial expressive role in explicating inferences (and therefore conceptual contents) that he assigns to it. These are picked up and emphasized, to some extent even in his own later work, but especially in the subsequent logistical tradition to which he gave birth, and make it difficult to work back to an appreciation of his original logical project. The difficulties stem from his use of the now-classical two-valued conditional.

It is difficult now to read the definition by which he introduced his conditional (in the fifth paragraph of the *Begriffsschrift*) without being blinded by hindsight—in particular by the glare of the truth-tabular tautology formulation presented by Wittgenstein in the *Tractatus*. Frege does define what has come to be called (by the lights of this work ludicrously inappropriately) the 'material' conditional. He does so, however, not in terms of a semantic distinction between judgeable contents that are true and those that are false, but rather in terms of a pragmatic distinction between those that are affirmed (*bejaht*) and those that are denied (*verneint*). This is his invariable practice in the *Begriffsschrift*, although in later years he is happy enough to recast these claims in terms of truth (as part of the reorientation of his thought toward logical truth that Dummett rightly complains about). Putting things in terms of truth rather than affirmation pushes into the background (though it does not abolish) the way in which the semantic notion of content is beholden to the pragmatic notion of force, in the explanation of which it serves, and which is the source of the priority of judgeable contents and so, even in the later work, of the special central and ineliminable role played by the True as *Bedeutung*. It is worth recalling in this connection Frege's formulation of his view in 1915, already quoted in Section II above: "'True' only makes an abortive attempt to indicate the essence of logic, since what logic is really concerned with is not contained in the word 'true' at all but in the assertoric force with which a sentence is uttered . . . the thing that indicates most clearly the essence of logic is the assertoric force with which a sentence is uttered."<sup>68</sup>

In fact Frege's view is that 'true' is a bit of logical vocabulary, which serves to express explicitly what is done implicitly in asserting. This is why by the time of the *Grundgesetze* (1893) he has adopted a regimentation in which all claims are expressed explicitly in the form of identities that have a sentence on one side and the canonical name 'the True' on the other. It is with that regimentation in mind, in turn, that he claims that the True is an object that must be recognized, at least implicitly, by anyone who makes judgments at all. (Identity claims are explained as the explicit expression of "recognition judgments" in the *Grundlagen*; see 7.1 below.) Assimilating all assertions to assertions of identities permits the use of his (ultimately substitutional) semantics for identity statements (forwarded in "Über Sinn und Bedeutung") in general application to all claims, which is his strategy in the *Grundgesetze*.

The key point is that explicitation is not explanation. Proprieties of inference are not *explained* in terms of something more primitive by being expressed in the explicit form of claims by the use of conditionals; the force of asserting or judging is not *explained* by expressing it explicitly as a saying of a sentence that it is (a name of the) true. This is why Frege always insists that truth is indefinable, something the understanding of which is always already implicit in claiming. "The True" is not a name whose sense one can grasp first, and then appeal to in explaining what it is to make a claim; its use merely makes explicit what is implicit in claiming. It has an expressive, not an explanatory role. Thinking of it the other way around is making a mistake with respect to 'true' and claiming that is strictly analogous to the inferential formalist's mistake regarding the conditional and inferring. A version of the preferred pragmatic direction of explanation is presented below, where Chapter 5 discusses the role of 'true' in terms of the expressive, explicitating function it performs with respect to claiming, according to the account of that practice offered in Chapter 3.

Bracketing subtleties regarding the relation between truth and commitment or affirmation, the fact remains that the conditional Frege actually defines and employs rules out only the case in which the consequent is denied or taken to be false while the antecedent is affirmed or taken to be true. This form of conditional, whatever its compositional virtues, is an extremely impoverished resource for the expression of proprieties of inference. The job that has been attributed to the conditional for Frege is that a conditional be affirmable or taken to be true just in case the inference from its antecedent to its consequent is endorsed or taken to be correct. Using the two-valued conditional to establish the connection between the correctness of an inference and the truth or endorsement of the claims that are its premises and conclusions has unpalatable results. Frege clearly has in mind a fundamental semantic principle regarding this connection: a good inference never takes one from premises that are true to a conclusion that is not true. This is a way of thinking about inferences as commitment-preserving: if one is committed to the premises of a good inference, in the sense of taking them



to be true (the sense that matters for assertion and judgment), then one is committed in the same sense to the conclusion.<sup>69</sup>

Such a principle could be agreed to both by those who adopt the traditional order of semantic explanation—by understanding the principle as explicating the correctness of inference in terms of a prior notion of truth (or taking-true)—and by those who adopt the converse order of semantic explanation (pursued in this work)—by taking truth or the sort of commitment involved in taking-true (and hence in asserting and judgment) to be explicated as what is preserved by good inferences.<sup>70</sup> But in either case, that truth or commitment is preserved by an inference ought to be taken to be a *necessary* condition of its being a good inference, not a *sufficient* condition. Affirming or taking-true both the claim that Hegel was Hölderlin's roommate and the claim that 43 is prime, and so being committed to the inference from the one claim to the other preserving truth and commitment, does not involve endorsing the propriety of that inference.

The two-valued conditional is subject to this familiar sort of complaint about fallacies of irrelevance precisely because the inferences it codifies explicitly are those that result from implausibly treating the plausible semantic preservation principle as, not only a necessary condition of good inference, but also as a sufficient one. It follows that the two-valued conditional Frege actually defines is an alarmingly bad choice for making explicit actual proprieties of inference. That fact in turn seems to cast doubt on the expressive understanding of his project. If he really wants logical vocabulary to make inferences explicit—because he wants to make nonlogical contents explicit and understands them in inferential terms—why does he employ the blunt, crude tool that is the two-valued conditional, whose expressive powers are hopelessly inadequate for the task of expressing the material inferences that might plausibly be identified with conceptual contents?

The answer is that although he hopes eventually to be able to use logical vocabulary to make explicit the inferential involvements in virtue of which nonlogical claims have the conceptual contents they do, the task Frege actually undertakes in the text of the *Begriffsschrift* is much less ambitious. The only concepts whose inferential role he actually makes explicit there are the logical concepts themselves, and those mathematical concepts that turn out to be definable from them. The concepts of geometry and mechanics—and indeed, the rest of the nonlogical concepts that philosophers might be interested in clarifying by expressing them explicitly—are to be *expressible* by means of logical vocabulary, together with other primitive signs. They are not understood to be, as some of the mathematical concepts (but not, for instance, those of geometry) are, *definable* by means of the logical vocabulary *alone*.

The first stage of Frege's grand project of clarification of nonlogical concepts through their explicitation in logical terms is to make explicit the conceptual contents of the logical expressions that are to be employed in that project. These concepts must themselves be "formed scientifically." This is

why he is proud to display, for each of the official propositions of the *Begriffsschrift* (couched entirely in logical vocabulary), what it follows from (in the proof of the proposition) and what follows from it (in the appendix). Doing so specifies the inferential role of those propositions, and so, indirectly, the conceptual content of the subsentential logical vocabulary that occurs in them.

Thus the only inferences Frege makes explicit in the *Begriffsschrift* are the inferences that are good in virtue of their logical form—for these determine the conceptual content (in his sense) of his logical vocabulary. He finds, in the two-valued conditional, an expressive equilibrium: the inferences in virtue of which that conditional means what it means can themselves be expressed and codified by the use of that conditional. Frege's logical vocabulary is potentially (and he makes it actually) self-explicating.<sup>71</sup> The official propositions of the *Begriffsschrift* explicitly specify the inferential roles of the logical vocabulary, and the inferential roles of those propositions can be expressed explicitly in terms of that vocabulary. Fascinated by how much of mathematical vocabulary turns out to be logical vocabulary in this sense, Frege does not in this work pursue the question of the expressive adequacy of his conditional for material, nonlogical, conceptual contents. He devotes most of the rest of his life to exploring the conceptual contents that can be made explicit by the use of this extensional conditional.<sup>72</sup>

The results he achieves with the poor expressive resources of the two-valued conditional deserve our awe and admiration. Nevertheless, the motivations remain for the grander semantic expressive aspirations that the young Frege contributes to the inferentialist tradition. It was pointed out above that distinguishing a privileged class of good inferences as good in virtue of their logical form, that is, as logically valid inferences, requires being able to pick out some vocabulary as distinctively *logical* vocabulary. Then the logically valid inferences are just those good inferences that remain good on all substitutions of nonlogical for nonlogical vocabulary. The demarcational question of how logical locutions ought to be identified has received various influential answers. The current suggestion is that Frege's early work is predicated on the idea that what distinguishes vocabulary as specifically *logical* is its expressive role in making conceptual content explicit. Vocabulary deserves the appellation 'logical' just in case it serves to make *explicit*, as the content of a claim, proprieties concerning the use of the expression that otherwise remain *implicit* in practice, specifically the proprieties in virtue of which it has the conceptual content that it does. It is because Frege understands those content-conferring practical proprieties to be in the first instance proprieties of *inferential* practice that the paradigmatic sentential logical locution for him is the conditional. One of the central tasks of the rest of this work is to show how this semantic expressive paradigm can be extended to other logical and semantic locutions.

In the next chapter it is argued that a key link connecting the implicit

norms governing the use of expressions with the conceptual content those practices confer on them is provided by the notion of the *incompatibility* of commitments. In practical terms of normative status, to treat  $p$  and  $q$  as incompatible claims is to take it that commitment to one precludes entitlement to the other. Practices properly articulated to be interpretable as instituting the normative statuses of commitment and entitlement required for incompatibility relations are thereby interpretable as conferring semantic content on the states, attitudes, and performances that stand in incompatibility relations. The content of a claim can be represented by the set of claims that are incompatible with it. For instance, a relation of entailment, required for an inferential semantics, can be derived according to the principle that  $p$  entails  $q$  just in case everything incompatible with  $q$  is incompatible with  $p$ . The formal semantics generated by such incompatibility interpretations is quite rich. It has been shown, for instance, how to represent classical logic, relevance logic, and various systems of orthologic (or quantum logic) by constraints on incompatibility relations.<sup>73</sup>

*Negation*, as a logical connective supporting formally valid inferences, plays the same explicating role with respect to material *incompatibility* relations among judgeable (that is propositional) contents that the *conditional* plays with respect to material *inferential* relations among such contents. The formal negation of a claim is constructed as its minimal incompatible, the claim that is entailed by each one of the claims incompatible with the claim of which it is the negation. Thus in the context of a conditional that makes entailment relations explicit, the introduction of a locution playing the inferential role of negation makes it possible to make explicit the relation of material incompatibility between claims. To assert that  $p$  is incompatible with  $q$ , one asserts the conditional whose antecedent is  $p$  and whose consequent is the negation of  $q$ . Conjunction and disjunction can be handled straightforwardly as corresponding to Boolean operations on the sets of incompatibles that represent conceptual contents according to this sort of semantic model.

Chapter 7 below discusses Frege's treatment of identity locutions as making explicit the substitution-inferential commitments that are implicit in the use of singular terms. It also shows how that idea can be extended to an account of the use of quantifiers as making explicit the different sort of substitution-inferential commitment that is implicit in the use of predicates. The job of the next chapter is to offer an account of the normative practices of claiming and judging, and of the propositional contents conferred on states, attitudes, performances, and expressions by their playing appropriate roles in those practices. This account gives a definite sense to the notion of explicit *sayings*, in terms of norms implicit in *doings*. What is explicit is then the propositional content that is said or believed. In this fundamental sense, H. L. Mencken makes the content of his thought explicit, and expresses it fully, by asserting the declarative sentence: "Natives of Appalachia are clay-eating

sub-humans." But it is also possible to use logical vocabulary to make explicit expressively essential inferential involvements that remain implicit in the concepts employed in making this claim.

In subsequent chapters various other locutions are introduced as being used so as to make explicit, in this sense, some feature of the practices that originally confer propositional content (so that having such contents can be understood to consist in how it is correct for those locutions to be used, according to the practices in question). Not only the standard logical vocabulary, but also traditional semantic vocabulary such as 'true', 'refers', and the 'of' of intentional aboutness, should be understood as semantically explicating. The point of using these sorts of expressions is to make explicit as the contents of claims (whose consequences can be explored and which can be justified and disputed) some critical element of the practices of talking and believing in virtue of which it is possible to interpret anything as propositionally contentful in the first place. Furthermore, another range of expressions, including such locutions as 'claims that', 'believes that', 'intends that', and normative talk of commitments and entitlements, is interpreted as *pragmatically* explicating. The point of using these sorts of expressions is to make explicit as the contents of claims some of the pragmatic elements of the practices of talking, believing, and acting that confer propositional contents. One thread running through the later chapters of this work is the attempt to achieve an analog of the expressive equilibrium Frege achieves in the propositional fragment of the *Begriffsschrift*. The challenge is to show how not only the semantics, but the pragmatics outlined in the first four chapters can be made explicit, in terms of vocabulary that is introduced by specifying practices of using it that are sufficient to confer on it the content that is then employed in making explicit precisely those practices and that content. The ideal is that the theory should specify practices sufficient to confer on the various locutions considered all the kinds of content required to state the theory itself.

## V. CIRCUMSTANCES AND CONSEQUENCES OF APPLICATION

### 1. Dummett's Model

The previous section of this chapter introduced three themes: that conceptual content is to be understood in terms of role in reasoning rather than exclusively in terms of representation, that the capacity for such reasoning is not to be identified exclusively with mastery of a logical calculus, and that besides theoretical and practical reasoning using contents constituted by their role in material inferences, there is a kind of *expressive* rationality that consists in making implicit, content-conferring inferential commitments explicit as the contents of assertional commitments. Being

rational in the primary sense is having states and attitudes and producing performances that have propositional contents. The next chapter discusses how the inferential articulation essential to such contents is conferred on them by the way in which the states, attitudes, and performances exhibiting those contents are caught up in the game of giving and asking for reasons. Rationality consists in mastery of those practices. It is not to be understood as a logical capacity. Rather, specifically logical capacities presuppose and are built upon underlying rational capacities. The fundamental characteristic role of logical vocabulary is to make it possible to talk and think *explicitly* about the inferentially articulated semantic contents *implicitly* conferred on expressions (among other things) by their role in rational practice. The optional introduction of sophisticated logical explicating vocabulary has an expressive point and payoff. By its means the material inferential practices, which govern and make possible the game of giving and asking for reasons, are brought into that game (and so into consciousness) as explicit topics of discussion and justification. In this way, in the context of the three basic themes mentioned above, an *expressive* understanding of logic was introduced—according to which formal validity of inferences is substitutionally defined in terms of material correctness of inferences together with the discrimination of some privileged vocabulary; that privileged vocabulary is identified as logical vocabulary; and what it is for something to be a bit of logical vocabulary is explained in terms of its semantically expressive role.

These ideas, to be found in the early works of Frege and Sellars, provide the beginnings of the structure within which modern inferentialism develops. The approach they suggest can be made more definite by considering a general model of conceptual contents as inferential roles that has been recommended (in somewhat different terms) by Dummett. According to that model, the use of any linguistic expression or concept has two aspects: the circumstances under which it is correctly applied, uttered, or used, and the appropriate consequences of its application, utterance, or use. Though Dummett does not put the point this way, this model connects to inferentialism of the Sellarsian sort via the principle that part of the content to which one is committed by using the concept or expression may be represented by the material inference one implicitly endorses by such use: the inference from the circumstances of appropriate employment to the appropriate consequences of such employment.

The original source for the model lies in a treatment of the grammatical category of sentential connectives. Dummett's two-aspect model is a generalization of a standard way of specifying the inferential roles of logical connectives, due ultimately to Gentzen. Gentzen defined connectives by specifying *introduction rules*, or inferentially sufficient conditions for the employment of the connective, and *elimination rules*, or inferentially necessary consequences of the employment of the connective. So, to define the

inferential role of an expression '&' of Boolean conjunction, one specifies that anyone who is committed to  $p$ , and committed to  $q$ , is thereby to count also as committed to  $p \& q$ , and that anyone who is committed to  $p \& q$  is thereby committed both to  $p$  and to  $q$ . The first schema specifies, by means of expressions that do not contain the connective, the *circumstances* under which one is committed to claims expressed by sentences that contain (as principal connective)<sup>74</sup> the connective whose inferential role is being defined, that is, the sets of premises that entail them. The second schema specifies, by means of expressions that do not contain the connective, the *consequences* of being committed to claims expressed by sentences that contain (as principal connective) the connective whose inferential role is being defined, that is, the sets of consequences that they entail.

Dummett makes a remarkable contribution to inferentialist approaches to conceptual content by showing how this model can be generalized from the case of logical connectives to provide a uniform treatment of the meanings of expressions of other important grammatical categories: sentences, predicates and common nouns, and singular terms. The application to the *propositional* contents expressed by whole declarative sentences is straightforward. What corresponds to an introduction rule for a propositional content is the set of inferentially sufficient conditions for asserting it, and what corresponds to an elimination rule is the set of inferentially necessary consequences of asserting it, that is, what follows from doing so. Dummett says: "Learning to use a statement of a given form involves, then, learning two things: the conditions under which one is justified in making the statement; and what constitutes acceptance of it, i.e., the consequences of accepting it."<sup>75</sup> Dummett presents his model as specifying two fundamental features of the use of linguistic expressions. In what follows, it is applied in the context of the previous ideas, to bring into relief the implicit material inferential content a concept or expression acquires in virtue of being used in the ways specified by these two "aspects." The link between *pragmatic significance* and *inferential content* is supplied by the fact that asserting a sentence is (among other things) implicitly undertaking a commitment to the correctness of the material inference from its circumstances to its consequences of application.

Dummett applies his model exclusively to conceptually contentful *linguistic* expressions. But it is clear that the model has a wider application—to intentional states and attitudes generally. Dummett's model just provides some structure to the representation of the functional roles of intentional states. For instance, one could think about the functional role played by a belief with a particular propositional content in terms of the circumstances in which it is appropriate to acquire a belief with that content, and the appropriate consequences of such acquisition. In the next two chapters this basic Dummettian structure is further articulated, to allow a richer representation of the functional roles of states, performances, and expressions.

## 2. *Inferential Connection between Even Noninferential Circumstances and Consequences of Application*

The concepts least easily assimilated to an inferential model are the empirical concepts whose core employment is in perception and the formulation of observation reports. For such reports are essentially *noninferential*—in the sense that they are elicited as responses to features of the largely nonlinguistic environment, rather than as conclusions drawn from other claims. Their content accordingly derives (at least in large part) from the *reliable differential responsive dispositions* that those who have mastered the concepts exhibit with respect to their application. Such concepts can be assimilated to the inferentialist understanding of conceptual contents by adapting Dummett's idea of distinguishing two crucial features of the use of linguistic expressions: their *circumstances* of appropriate application, and the appropriate *consequences* of such application. In terms of this model, it is possible to understand the use of *any* expression as implicitly involving an inferential commitment. In particular, by using the expression, one is (among other things) committed to the propriety of the inference from its circumstances to its consequences of application. The consequences of application are always themselves inferentially related to the concept in question (although the inference involved may include practical inferences, whose conclusions are commitments to act). The circumstances of application need not themselves be linguistic. For the concept *red*, for instance, they include the presence of visibly red things. Nonetheless, even the use of concepts of this sort can be seen to embody inferential commitments, to the propriety of applying inferential consequences of *red*—for instance, *colored*—to anything that *red* is properly applied to.<sup>76</sup>

It is in this way that a broadly inferential approach can incorporate into its conception of the contents of empirical concepts the nonlinguistic circumstances in which they are correctly noninferentially applied. Thus the concepts *water* and its twin-earth analog *twater*, which are by hypothesis alike except that one is appropriately applied to H<sub>2</sub>O and the other to XYZ, count as involving different *inferential* contents. This is so even though they fund inferential moves involving the same sorts of *noises*, from saying "That's water" to "That's liquid," for instance. For they involve different circumstances of appropriate application, and hence different inferential transitions from those circumstances to their consequences. So even though it is the practices of those whose concepts they are that confer on them their contents, the earthlings and twin-earthlings need not be able to tell that they have different concepts, if water and *twater* are indistinguishable to them. They are not omniscient about the inferential commitments implicit in their own concepts. For the interpreter who is making sense of their practices, and who is able (not perceptually, but conceptually) to distinguish H<sub>2</sub>O and XYZ, can understand transported earthlings as *mistaking* the XYZ they look at for

water, as *inappropriately* applying the concept they express with their word 'water' to that unearthly stuff. The circumstances of appropriate noninferential application of the concept expressed by the English word 'water' require that it be applied in response to a sample of  $H_2O$ .

In this way the circumstances of appropriate application of a claim can include not only other claims (from which the one in question could be inferred) but also perceptual circumstances (to which one has been trained to respond *noninferentially* by endorsing the target claim). The appropriate consequences of application of a claim can include not only the inferential acquisition of further beliefs whose contents follow from the contents of the belief in question but also, in the context of further contentful intentional states, the noninferential responsive performance of *actions*, under the descriptions by which they can be exhibited as the conclusions of practical inferences. This is explicitly acknowledged in the continuation of the passage from Dummett quoted above: "Here 'consequences' must be taken to include both the inferential powers of the statement and anything that counts as acting on the truth of the statement." So the circumstances and consequences of application Dummett is talking about should not be identified with inferentially necessary and sufficient conditions, where this means identifying them with sets of claims or beliefs that are conclusions or premises of theoretical inferences involving the content in question. It turns out, however, that the circumstances and consequences model can be understood as an inferential model, regardless of whether the circumstances and consequences are themselves already thought of in inferential terms. The inferential element enters this picture in the commitment undertaken by one who employs a given content to the propriety of the transition from the circumstances of appropriate application to the appropriate consequences of application of a conceptual content. This will be construed as a broadly inferential commitment, though the detailed justification for this characterization will not emerge until the next chapter.<sup>77</sup>

One advantage of thinking about conceptual content as determining functional role specified in terms of proprieties governing circumstances and consequences of application is the room it makes for a *pragmatic* picture of understanding or grasping such a content. Understanding can be understood, not as the turning on of a Cartesian light, but as practical mastery of a certain kind of inferentially articulated *doing*: responding differentially according to the circumstances of proper application of a concept, and distinguishing the proper inferential consequences of such application. This is not an all-or-none affair; metallurgists understand the concept *tellurium* better than most of us do, for training has made them master of the inferential intricacies of its employment in a way that we can only crudely approximate. On this inferentialist rendering, thinking clearly is a matter of knowing what one is committing oneself to by a certain claim, and what would entitle one to that commitment. Writing clearly is providing enough clues for a reader to infer



what one intends to be committed to by each claim, and what one takes it would entitle one to that commitment. Failure to grasp either of these components is failure to grasp the inferential commitment that use of the concept involves, and so failure to grasp its conceptual content.

### 3. *One-Sided Theories of Meaning*

Verificationists, assertibilists, and reliabilists make the mistake of treating the first aspect as exhausting content. Understanding or grasping a content is taken to consist in practically mastering the circumstances under which one becomes entitled or committed to endorse a claim or acquire a belief, quite apart from any grasp of what one becomes entitled or committed to by such endorsement or acquisition. But claims can have the same circumstances of application and different consequences of application, as for instance "I foresee (or predict) that I will write a book about Hegel" and "I will write a book about Hegel" do. Any circumstances under which one is entitled to one of these claims (or to acquire the belief it expresses) are circumstances under which one is entitled to the other. (If this does not seem right for the actual concepts expressed by 'foresee' and 'predict', artificial variants clearly can be constructed for which it is.) Yet what follows from the claims is quite different. If I will write a book about Hegel, then I will write a book about Hegel. Yet that I will write a book about Hegel does not follow from my foreseeing or predicting that I will (as the sad history of orphaned first volumes of ambitious projects attests). The consequences of these claims are quite different. Examples meeting the conditions required for this point are forthcoming in any idiom expressively rich enough to contain *pragmatically* explicitating locutions, which permit one to *say* what one is *doing* in performing a certain speech act or acquiring a certain state or attitude—for instance "I claim that *p*" or "I believe that *p*." (These locutions are discussed in Chapter 8.)

In any idiom expressively rich enough to contain *semantically* explicitating locutions, whose paradigm is the conditional, the difference in inferential consequences of application between the sentence whose utterance performs a speech act (a doing in which the force is left implicit) and the sentence whose utterance explicitly *says* that that is what one is doing (so that force becomes part of the content) itself becomes explicit in the use of conditionals with those sentences as antecedents. The circumstances of appropriate application or assertibility conditions of the conditionals "If I will write a book about Hegel, then I will write a book about Hegel" and "If I foresee (or predict) that I will write a book about Hegel, then I will write a book about Hegel" are quite different. The assertibility of the second conditional, but not the first, depends on auxiliary hypotheses about how good at foreseeing or predicting I am. So for idioms that contain both the pragmatically explicitating locutions that permit the construction of pairs of sentences with identical

circumstances of application and different consequences of application, and the semantically explicating locutions that permit the construction of conditionals whose circumstances of application differ depending on the consequences of application of their antecedents, it is possible to show the inadequacy of a semantics that avails itself only of assertibility conditions or circumstances of appropriate application. For examples of the sort just considered show that substituting another sentence with the same assertibility conditions for a sentence that is the antecedent of a conditional can alter the assertibility conditions of the compound. In such an expressively rich environment, then, assertibility conditions cannot provide an adequate model of what Dummett calls "*ingredient content*," the contribution the occurrence of a sentence makes to the use of sentences in which it appears as a component. But this fact simply reflects the inadequacy of the model for the expression of conceptual content as inferential role, even in the more expressively impoverished idioms in which the pragmatically and semantically explicating locutions are not available.<sup>78</sup>

The inadequacy of a notion of semantic content that is restricted to circumstances of application to the exclusion of consequences of application has already appeared in another guise above. The point of the discussion of Sellars's application of inferentialist ideas to the understanding of noninferential reports, in Section III, was that parrots and photocells and so on might reliably discriminate the circumstances in which the concept *red* should be applied, without thereby grasping that concept. This would happen precisely in the case where they have no mastery of the consequences of such application—when they cannot tell that it follows from something being red that it is colored, that it is not a prime number, and so on. You do not convey to me the content of the concept *gleeb* by supplying me with an infallible gleebsness tester that lights up when and only when exposed to gleeb things. I would in that case know what things were gleeb, without knowing what I was saying about them when I called them that, what I had found out about them or committed myself to. Dummett offers two examples of philosophically important concepts where it is useful to be reminded of this point: "An account, however accurate, of the conditions under which some predicate is rightly applied may thus miss important intuitive features of its meaning; in particular, it may leave out what we take to be the point of our use of the predicate. A philosophical account of the notion of truth can thus not necessarily be attained by a definition of the predicate 'true', even if one is possible, since such a definition may be correct only in the sense that it specifies correctly the application of the predicate, while leaving the connections between this predicate and other notions quite obscure."<sup>79</sup> Even more clearly: "A good example would be the word 'valid' as applied to various forms of argument. We might reckon the syntactic characterization of validity as giving the criterion for applying the predicate 'valid' to an argument, and the semantic characterization of validity as giving the consequences of such an application . . . If [a student] is taught in a very unimaginative way,

he may see the classification of arguments into valid and invalid ones as resembling the classification of poems into sonnets and non-sonnets, and so fail to grasp that the fact that an argument is valid provides any grounds for accepting the conclusion if one accepts the premises. We should naturally say that he had missed the point of the distinction."

Pragmatists of the classical sort, in contrast, make the converse mistake of identifying propositional contents exclusively with the consequences of endorsing a claim: looking downstream to the claim's role as a premise in practical reasoning and ignoring its proper antecedents upstream. The fact that the pragmatist's emphasis is on practical consequences is not relevant to this complaint. The problem is that one can know what follows from the claim that someone is responsible for a particular action, that an action is immoral or sinful, that a remark is true or in bad taste, without for that reason counting as understanding the claims involved, if one has no idea when it is appropriate to make those claims or apply those concepts. Being classified as AWOL does have the consequence that one is liable to be arrested, but the specific circumstances under which one acquires that liability are equally essential to the concept.

It was pointed out that Frege's practice in the *Begriffsschrift* is to specify both the circumstances and the consequences of application of his claims, which in the context of that project (excluding as it does concepts with empirical and practical content deriving from their relation to perception and action) can be identified with the inferentially sufficient premises from which they follow and the inferentially necessary conclusions they lead to. Yet his official definition of conceptual content refers only to consequences, and Carnap follows him in this regard. For the special sort of concepts they are concerned with, where only inferential circumstances and consequences are in play, this restriction does not amount to ignoring circumstances of application. Restricting consideration for the sake of an example to one-premise inferences, associating with each sentence the set of sentences that follow from it determines for each sentence which sentences it follows from. So at the global level, nothing is lost by officially defining content in terms of inferential consequences alone. As will emerge below, it is quite otherwise when one is concerned locally with the content associated with each sentence—for instance in asking what it is to understand the content expressed by one sentence (but perhaps not another), or to alter the content expressed by one sentence, or to introduce a new content. Nor will the technical dodge of restriction of content to consequences be adequate when attention is turned to the sort of empirical and practical content concepts get from their involvement in perception and action.

#### 4. *Conservativeness and the Coherence of Logical Concepts*

Of course, such one-component theories do not simply ignore the aspects of content they do not treat as central. Dummett says:

Most philosophical observations about meaning embody a claim to perceive . . . a simple pattern: the meaning of a sentence consists in the conditions for its truth and falsity, or in the method of its verification, or in the practical consequences of accepting it. Such dicta cannot be taken to be so naive as to involve overlooking the fact that there are many other features of the use of a sentence than the one singled out as being that in which its meaning consists: rather, the hope is that we shall be able to give an account of the connection that exists between the different aspects of meaning. One particular aspect will be taken as central, as constitutive of the meaning of any given sentence . . . ; all other features of the use of the sentence will then be explained by a uniform account of their derivation from that feature taken as central.<sup>80</sup>

Pursuing this notion of derivation provides a helpful perspective on the idea of conceptual contents articulated according to material inferences, and on the role of explicit inference licenses such as conditional statements in expressing and elucidating such inferences and so such contents.

The strategy of attempting to *derive* one aspect of the use of an expression (or the significance of an intentional state) from another—in particular to derive appropriate consequences of application from circumstances of appropriate application, or vice versa—expresses Dummett's appreciation of the need for the semantic theorist to be able to explain two crucial features of our practices regarding conceptual contents. Concept-users are often confronted with decisions regarding *alternative* concepts and so are obliged to decide not only that certain uses of a given concept should be rejected as incorrect but also that certain concepts should themselves be rejected as inadequate or incorrect. We criticize our concepts and sometimes reject them. Furthermore, doing so is not simply a matter of free or arbitrary stipulation. Criticism of our concepts is constrained and sometimes compelled. These are important phenomena—an attempt to take proper account of them guides the discussion below. Dummett acknowledges them as motivating the theoretical acknowledgment of a need for *harmony* between the circumstances and consequences of application: "A naive view of language regards assertibility-conditions for a statement as exhausting its meaning: the result is to make it impossible to see how a meaning can ever be criticized, revised, or rejected; it was just such a naive view which led to the use of the notorious 'paradigm-case argument'. An almost equally naive view is that which distinguishes the assertibility-conditions of a statement as its 'descriptive meaning' and its consequences as its 'evaluative meaning', dispensing with any requirement of harmony between them, but holding that we have the right to attach whatever evaluative meaning we choose to a form of statement irrespective of its descriptive meaning."<sup>81</sup>

For the special case of defining the inferential roles of logical connectives by pairs of sets of rules for their introduction and for their elimination, which

motivates Dummett's broader model, there is a special condition it is appropriate to impose on the relation between the two sorts of rules. "In the case of a logical constant, we may regard the introduction rules governing it as giving conditions for the assertion of a statement of which it is the main operator, and the elimination rules as giving the consequences of such a statement: the demand for harmony between them is then expressible as the requirement that the addition of the constant to a language produces a conservative extension of that language."<sup>82</sup> Recognition of the appropriateness of such a requirement arises from consideration of connectives with "inconsistent" contents. As Prior pointed out, if a connective, which after Belnap may be called 'tonk', is defined as having the introduction rule proper to disjunction and the elimination rule proper to conjunction, then the first rule licenses the transition from  $p$  to  $p \text{ tonk } q$ , for arbitrary  $q$ , and the second licenses the transition from  $p \text{ tonk } q$  to  $q$ . The result is what he called a "runabout inference ticket," which permits any arbitrary inference.

Prior thought that this possibility shows the bankruptcy of Gentzen-style definitions of inferential roles. Belnap shows rather that when logical vocabulary is being introduced, one must constrain such definitions by the condition that the rule not license any inferences involving only old vocabulary that were not already licensed before the logical vocabulary was introduced.<sup>83</sup> That is, it must be ensured that the new rules provide an *inferentially conservative extension* of the original field of inferences. From the point of view of the joint commitments to understanding conceptual content in terms of material inference and conceiving the distinctive role of logical vocabulary as making those content-conferring inferential connections explicit in the form of claims, this constraint on the definition of logical particles by introduction and elimination rules makes perfect sense. For if those rules are not inferentially conservative, the introduction of the new vocabulary licenses new material inferences and so alters the contents associated with the old vocabulary. The expressive approach to logic motivates a criterion of adequacy for introducing logical vocabulary to the effect that no *new* inferences involving only the *old* vocabulary be made appropriate thereby. Only in this way can logical vocabulary play the expressive role of making explicit the original material inferences and so nonlogical conceptual contents.

### 5. *Nonlogical Concepts Can Incorporate Materially Bad Inferences*

The problem of what Dummett calls a lack of "harmony" between the circumstances and the consequences of application of a concept can arise, however, not only for logical vocabulary but also for concepts with material contents. Seeing how it does provides further help in understanding the notion of expressive rationality and the way in which the explicating role of logical vocabulary contributes to the clarification of concepts. For conceptual change can be:

motivated by the desire to attain or preserve a harmony between the two aspects of an expression's meaning. A simple case would be that of a pejorative term, e.g. 'Boche'. The condition for applying the term to someone is that he is of German nationality; the consequences of its application are that he is barbarous and more prone to cruelty than other Europeans. We should envisage the connections in both directions as sufficiently tight as to be involved in the very meaning of the word: neither could be severed without altering its meaning. Someone who rejects the word does so because he does not want to permit a transition from the grounds for applying the term to the consequences of doing so. The addition of the term 'Boche' to a language which did not previously contain it would produce a non-conservative extension, i.e. one in which certain other statements which did not contain the term were inferable from other statements not containing it which were not previously inferable.<sup>84</sup>

This crucial passage makes a number of points that are worth untangling.

First of all, it shows how concepts can be criticized on the basis of substantive beliefs. If one does not believe that the inference from German nationality to cruelty is a good one, one must eschew the concept *Boche*. For one cannot deny that there are any Boche—that is just denying that anyone is German, which is patently false. One cannot admit that there are Boche and deny that they are cruel—that is just attempting to take back with one claim what one has committed oneself to with another. One can only refuse to employ the concept, on the grounds that it embodies an inference one does not endorse. (When the prosecutor at Oscar Wilde's trial asked him to say under oath whether a particular passage in one of his works did or did not constitute blasphemy, Wilde replied, "Blasphemy is not one of my words."<sup>85</sup>) Highly charged words like 'nigger', 'whore', 'Republican', and 'Christian' have seemed a special case to some because they couple "descriptive" circumstances of application to "evaluative" consequences. But this is not the only sort of expression embodying inferences that requires close scrutiny. The use of *any* concept or expression involves commitment to an inference from its grounds to its consequences of application. Critical thinkers, or merely fastidious ones, must examine their idioms to be sure that they are prepared to endorse and so defend the appropriateness of the material inferential transitions implicit in the concepts they employ. In Reason's fight against thought debased by prejudice and propaganda, the first rule is that material inferential commitments that are potentially controversial should be made explicit as claims, exposing them both as vulnerable to reasoned challenge and as in need of reasoned defense.

It is in this process that formal logical vocabulary such as the conditional plays its explicating role. It permits the formulation, as explicit claims, of the inferential commitments that otherwise remain implicit and unexam-

ined in the contents of material concepts. Logical locutions make it possible to display the relevant grounds and consequences and to assert their inferential relation. Formulating as an explicit claim the inferential commitment implicit in the content brings it out into the open as liable to challenges and demands for justification, just as with any assertion. In this way explicit expression plays an elucidating role, functioning to groom and improve our inferential commitments and so our conceptual contents—a role, in short, in the practices of reflective rationality that Sellars talks about under the heading of “Socratic method.”

But if Dummett is suggesting that what is wrong with the concept *Boche* is that its addition represents a nonconservative extension of the rest of the language, he is mistaken. Its nonconservativeness just shows that it has a substantive content, in that it implicitly involves a material inference that is not already implicit in the contents of other concepts being employed. This is no bad thing. Conceptual progress in science often consists in introducing just such novel contents. The concept *temperature* was introduced with certain criteria or circumstances of appropriate application and with certain consequences of application. As new ways of measuring temperature are introduced, and new consequences of temperature measurements adopted, the complex inferential commitment that determines the significance of using the concept of temperature evolves.

The proper question to ask in evaluating the introduction and evolution of a concept is not whether the inference embodied is one that is already endorsed (so that no new content is really involved) but whether that inference is one that ought to be endorsed. The problem with ‘Boche’ or ‘nigger’ is not that once we explicitly confront the material inferential commitment that gives them their content, it turns out to be *novel*, but that it can then be seen to be indefensible and inappropriate. We want to be aware of the inferential commitments our concepts involve, to be able to make them explicit, and to be able to justify them. But there are other ways of justifying them than showing that we were already implicitly committed to them, before introducing or altering the concept in question. Making implicit commitments explicit is only a necessary condition of justifying them.

Even in the cases where it does make sense to identify harmony of circumstances and consequences with inferential conservativeness, the attribution of conservativeness is always relative to a background set of material inferential practices, the ones that are conservatively extended by the vocabulary in question. Conservativeness is a property of the conceptual content only in the context of other contents, not something it has by itself. There can be pairs of logical connectives, either of which is all right by itself, but both of which cannot be included in a consistent system. It is a peculiar ideal of harmony that would be realized by a system of conceptual contents such that the material inferences implicit in every subset of concepts represented a conservative extension of the remaining concepts, in that no infer-

ences involving only the remaining ones are licensed that are not licensed already by the contents associated just with those remaining concepts. Such a system is an idealization because all of its concepts would already be out in the open, with none remaining hidden, to be revealed only by drawing conclusions from premises that have never been conjoined before, following out unexplored lines of reasoning, drawing consequences one was not previously aware one would be entitled or committed to by some set of premises. In short, this would be a case where Socratic reflection—making implicit commitments explicit and examining their consequences and possible justifications—would never motivate one to alter contents or commitments. Such complete transparency of commitment and entitlement is in some sense an ideal projected by the sort of Socratic practice that finds current contents and commitments wanting by confronting them with each other, pointing out inferential features of each of which we were unaware. But as Wittgenstein teaches in general, it should not be assumed that our scheme is like this, or depends upon an underlying set of contents like this, just because we are obliged to remove any particular ways in which we discover it to fall short.

These are reasons to part company with the suggestion, forwarded in the passage above, that inferential conservatism is a *necessary* condition of a “harmonious” concept—one that won’t “tonk up” a conceptual scheme. In a footnote, Dummett explicitly denies that conservativeness can in general be treated as a *sufficient* condition of harmony: “This is not to say that the character of the harmony demanded is always easy to explain, or that it can always be accounted for in terms of the notion of a conservative extension . . . The most difficult case is probably the vexed problem of personal identity.”<sup>86</sup> In another place, this remark about personal identity is laid out in more detail:

We have reasonably sharp criteria which we apply in ordinary cases for deciding questions of personal identity: and there are also fairly clear consequences attaching to the settlement of such a question one way or the other, namely those relating to ascriptions of responsibility, both moral and legal, to the rights and obligations which a person has . . . What is much harder is to give an account of the connection between the criteria for the truth of a statement of personal identity and the consequences of accepting it. We can easily imagine people who use different criteria from ours . . . Precisely what would make the criteria they used criteria for *personal identity* would lie in their attaching the same consequence, in regard to responsibility, motivation, etc., to their statements of personal identity as we do to ours. If there existed a clear method for deriving, as it were, the consequences of a statement from the criteria for its truth, then the difference between such people and ourselves would have the character of a factual disagreement, and one



side would be able to show the other to be wrong. If there were no connection between truth-grounds and consequences, then the disagreement between us would lie merely in a preference for different concepts, and there would be no right or wrong in the matter at all.<sup>87</sup>

Dummett thinks that there is a general problem concerning the way in which the circumstances and consequences of application of expressions or concepts ought to fit together. Some sort of "harmony" seems to be required between these two aspects of the use. The puzzling thing, he seems to be saying, is that the harmony required cannot happily be assimilated either to compulsion by facts or to the dictates of freely chosen meanings. But the options—matter of fact or relation of ideas, expression of commitment as belief or expression of commitment as meaning—are not ones that readers of "Two Dogmas of Empiricism" and its heirs ought to be tempted to treat as exhaustive.<sup>88</sup>

As already pointed out, talk of derivability is strictly stronger than talk of conservativeness. On the other side of the divide, the notion of a completely factual issue that Dummett appeals to in this passage is one in which the applicability of a concept is settled straightforwardly by the application of other concepts: the concepts that specify the necessary and sufficient conditions that determine the truth conditions of claims involving the original concept.<sup>89</sup> This conception, envisaged by a model of conceptual content as necessary and sufficient conditions, seems to require a conceptual scheme that is ideally transparent in the way mentioned above, in that it is immune to Socratic criticism. For that conception insists that these *coincide*—in that the individually sufficient conditions *already* entail the jointly necessary ones. Only then is it attractive to talk about content as truth conditions, rather than focusing on the substantive inferential commitments that relate the sufficient to the distinct necessary conditions, as recommended here. By contrast to the either/or that Dummett presents, in a picture according to which conceptual contents are conferred by being caught up in a social practical structure of inferentially articulated commitments and entitlements, material inferential commitments are a necessary part of any package of practices that includes material assertional or doxastic commitments. From this point of view, rendering conceptual content as truth conditions and thinking of them as necessary and sufficient conditions leaves out precisely the material content of concepts.

For the circumstances and consequences of application of a concept may stand in a substantive material-inferential relation. To ask what sort of "harmony" they should exhibit is to ask what material inferences we ought to endorse, and so what conceptual contents we ought to employ. This is not the sort of a question to which one ought to expect or even welcome a general or wholesale answer. Grooming our concepts and material inferential commitments in the light of our assertional commitments (including those we

find ourselves with noninferentially through observation) and the latter in the light of the former is a messy, retail business. Dummett thinks that a theory of meaning should take the form of an account of the nature of the "harmony" that ought to obtain between the circumstances and the consequences of application of the concepts we ought to employ. The present point is that one should not expect a theory of that sort to take the form of a specification of necessary and sufficient conditions for the circumstances and consequences of application of a concept to be harmonious. Rather, insofar as the idea of such a theory makes sense at all, it must take the form of an investigation of the ongoing elucidative process, of the "Socratic method" of discovering and repairing discordant concepts, which alone gives the notion of harmony any content. It is given content only by the process of *harmonizing* commitments, from which it is abstracted.

In Sellars's characterization of induction, introduced above, inductive inference is assigned an expressive role insofar as its conclusion is understood as being an inference license making explicit a commitment that is implicit in the use of conceptual contents antecedently in play. Rules of this sort assert an authority over future practice and for their entitlement answer both to the prior practice being codified and to concomitant inferential and doxastic commitments. In this way they may be likened to the principles formulated by judges at common law, intended both to codify prior practice, as represented by precedent, expressing explicitly as a rule what was implicit therein, and to have regulative authority for subsequent practice. The expressive task of making material inferential commitments explicit plays an essential role in the reflectively rational Socratic practice of harmonizing our commitments. For a commitment to become explicit is for it to be thrown into the game of giving and asking for reasons as something whose justification, in terms of other commitments and entitlements, is liable to question. Any theory of the sort of inferential *harmony* of commitments we are aiming at by engaging in this reflective, rational process must derive its credentials from its expressive adequacy to that practice, before it should be accorded any authority over it.

## 6. *Varieties of Inferentialism*

Section IV of this chapter introduced three related ideas:

1. the inferential understanding of conceptual content,
2. the idea of materially good inferences, and
3. the idea of expressive rationality.

These contrast, respectively, with

- 1'. an understanding of content exclusively according to the model of the representation of states of affairs,

- 2'. an understanding of the goodness of inference exclusively on the model of formal validity, and
- 3'. an understanding of rationality exclusively on the model of instrumental or means-end reasoning.

In this section these ideas were considered in relation to the representation of inferential role suggested by Dummett, in terms of the circumstances of appropriate application of an expression or concept and the appropriate consequences of such application. Both sections sought to introduce an *expressive* view of the characteristic role played by logical vocabulary and to indicate its relation to the practices constitutive of rationality.

One of the important benefits afforded by the emphasis in this section on understanding the Dummettian model of the use of linguistic expressions in terms of appropriate circumstances and consequences of application linked by an inferential commitment is the clarification it offers concerning the options that are open in working out an inferentialist approach to semantics. There are three different ways in which one might take inference to be of particular significance for understanding conceptual content. The weak inferentialist thesis is that inferential articulation is *necessary* for specifically *conceptual* contentfulness. The strong inferentialist thesis is that broadly inferential articulation is *sufficient* for specifically conceptual contentfulness—that is, that there is nothing more to conceptual content than its broadly inferential articulation. Dummett's model is particularly helpful for focusing attention on how important the qualification 'broadly' is in this formulation. For strong inferentialism as it is worked out in the rest of this project is *not* committed to the hyperinferentialist thesis, which maintains that *narrowly* inferential articulation is sufficient for conceptual contentfulness of all sorts.

The difference between the broad and the narrow conception of inferential articulation has three dimensions. First, and most important, the broad conception includes the possibility of *noninferential* circumstances and consequences of application. In this way (discussed in Chapter 4) the specifically *empirical* conceptual content that concepts exhibit in virtue of their connection to language entries in *perception* and the specifically *practical* conceptual content that concepts exhibit in virtue of their connection to language exits in *action* are incorporated into the inferentialist picture. The use of concepts with contents of these sorts can still be understood in terms of the material inferential commitment one who uses them undertakes: the commitment to the propriety or correctness of the inference from their circumstances to their consequences of application. Conceiving such inferences broadly means conceiving them as involving those circumstances and consequences, as well as the connection between them. The hyperinferentialist about conceptual content (adopting a position not endorsed here) would allow only *inferential* circumstances and consequences of application. Under

such a restriction, it is impossible to reconstruct the contents of actual concepts, except perhaps in some regions of mathematics.

Second, relations of *incompatibility* among claims and (so) concepts are considered broadly inferential relations, on grounds of their antecedents and their consequences. On the side of consequences, incompatibility relations underwrite the modal inferences codified by strict implication. For  $p$  entails  $q$  in this sense just in case everything incompatible with  $q$  is incompatible with  $p$ . So being a square entails being a rectangle, because everything incompatible with being a rectangle is incompatible with being a square. On the side of antecedents, the semantic relation of incompatibility will be understood (in the next chapter) in terms of the very same normative statuses of doxastic commitment and entitlement to such commitments, in terms of which inferences are construed (with commitment-preserving inferences corresponding roughly to deductive inferences, and entitlement-preserving inferences corresponding roughly to inductive inferences).

Finally, the notion of broadly inferential articulation is extended in subsequent chapters to include the crucial inferential substructures of *substitution* and *anaphora*. Substitutional commitments are defined as a species of inferential commitments (in Chapter 6) by distinguishing a class of substitution inferences. In this way the inferentialist paradigm can be extended so as to apply to the conceptual contents of subsentential expressions such as singular terms and predicates. Then anaphoric commitments are defined in terms of the *inheritance* of substitution-inferential commitments (in Chapter 7). In this way the inferentialist paradigm can be extended so as to apply to unrepeatable or token-reflexive expressions, such as demonstratives, indexicals, and pronouns.

It is important to keep in mind in reading what follows that the inferentialist project pursued here is a defense of the strong, not only the weak, inferentialist thesis. But it is not a form of hyperinferentialism. And while it eschews representational semantic primitives in favor of others more easily grounded in pragmatics, this is not because of a denial of the importance of the representational dimension of discursive practice. On the contrary, that choice serves rather an aspiration to make intelligible in a new way just what that representational dimension consists in.

## VI. CONCLUSION

### 1. *Grounding an Inferential Semantics on a Normative Pragmatics*

Inferentialism about conceptual content is not an explanatory strategy that can be pursued in complete abstraction from pragmatism about the norms implicit in the practical application of concepts. The considerations assembled here to motivate and recommend an inferentialist order of

semantic explanation appeal to a notion of *materially correct inferences*. In this chapter material proprieties of inference have been treated as primitives, playing the role of unexplained explainers. A critical criterion of adequacy by which such an approach should be assessed is clearly the extent to which a philosophically satisfying story can be told about these primitive proprieties of nonformal inference. The semantic theorist's entitlement to explanatory use of such primitives must be vindicated by situating the project of semantic theory in a broader context. Conceptual contents, paradigmatically propositional ones, are associated with linguistic expressions as part of an attempt to specify, systematically and explicitly, the correct *use* of those expressions. Such contents are associated with intentional states such as belief as part of a corresponding attempt to specify their behavioral significance—the difference those states make to what it is appropriate for the one to whom they are attributed to *do*.

The study of the practical significance of intentional states, attitudes, and performances (including speech acts) is *pragmatics*, as that term is used here. The projects of semantic theory and of pragmatic theory are intricately interrelated. If the semantic content and pragmatic context of a linguistic performance of a particular kind (paradigmatically assertion) are specified, a general theory of speech acts seeks to determine in a systematic way the pragmatic significance of that contentful performance in that context. But besides the direction of explanation involved in the local determination of pragmatic significance by semantic content, there is also a converse direction of explanation involved in the global conferral of semantic content by pragmatic significance. It must be explained how expressions can be used so as to confer on them the contents they have—what functional role the states they manifest must play in practice for them to be correctly interpreted as having certain intentional contents. Such an explanation amounts to an account of what it is for a state, attitude, performance, or expression to be *propositionally* contentful. Once a general notion of content has been made sense of in this way, particular attributions of contentfulness can then be offered as part of explanations or explicit specifications of the pragmatic significance of a state, attitude, performance, or expression.

The discussion of the next chapter should begin to make clearer just how a story about the conferral of content by practice is envisaged as relating to the use of attributions of content in the determination of pragmatic significances. One aspect of the situation of the semantic concept of content in a wider pragmatic context, however, is of particular relevance to the issue of entitlement to appeal to material proprieties of inference as semantic primitives. For the inferential proprieties that from the point of view of semantic theory are treated as primitive can be explained in the pragmatic theory as implicit in discursive practice (which includes intentional agency). An inferential move's normative status as correct or incorrect can be construed as instituted in the first instance by practical attitudes of taking or treating it

as correct or incorrect. The inferential norms that govern the use of expressions (or the significance of states, attitudes, and performances) are then understood as instituted by practical attitudes toward what the content is attributed to; they in turn confer that content on what it is attributed to.

Expressions come to mean what they mean by being used as they are in practice, and intentional states and attitudes have the contents they do in virtue of the role they play in the behavioral economy of those to whom they are attributed. Content is understood in terms of proprieties of inference, and those are understood in terms of the norm-instituting attitudes of taking or treating moves as appropriate or inappropriate in practice. A theoretical route is accordingly made available from what people *do* to what they *mean*, from their *practice* to the *contents* of their states and expressions. In this way a suitable pragmatic theory can ground an inferentialist semantic theory; its explanations of what it is in practice to treat inferences as correct are what ultimately license appeal to material proprieties of inference, which can then function as semantic primitives.

Sketching the possibility of such an explanatory path from attributions of practical attitudes to attributions of semantic content should help alleviate one sort of worry that might be elicited by the inferentialist invocation of materially correct inferences in explaining conceptual contentfulness. For otherwise the employment of a notion of *material* proprieties of inference in explaining *content* might seem blatantly circular. After all, are not materially good inferences just those that are good in virtue of the contents of the nonlogical concepts applied in their premises and conclusions, by contrast to the logically valid inferences, which are good in virtue of the logical form of those premises and conclusions? Presystematically, this is indeed how they should be thought of. But officially, the strategy is to start with proprieties of inference and to elucidate the notion of conceptual content in terms of those proprieties.

Talk of materially correct inferences is indeed intended to enforce a contrast with those that are formally correct (in the sense of logically valid). But the force of this contrast is just that the validity of inferences in virtue of their logical form is to be understood as a sophisticated, late-coming sort of propriety of inference, founded and conceptually parasitic on a more primitive sort of propriety of inference. This is the repudiation of the formalist approach to inference, for which the correctness of inference is intelligible only as formal logical validity, correctness in virtue of logical form. Calling the more primitive sort of propriety of inference *materially* correct simply registers the rejection of this order of explanation. It does not involve commitment to a prior notion of nonlogical content. If what it means to call an inference correct in the relevant sense can be explained without appeal to the use of logical concepts—for instance in terms of conduct interpretable as a practical taking or treating of an inference as correct—then there need be no circularity in appeal to such inferential proprieties in elaborating a notion of conceptual content.

2. *Knowing-That in Terms of Knowing-How, Formal Proprieties of Inference in Terms of Material Ones, Representational Content in Terms of Inferential Content*

A story that begins with inferring as a kind of practical *doing* and that leads to an account of the specifically *propositional* contentfulness of speech acts and intentional states holds out the promise of yielding an account of propositionally *explicit* saying, judging, or knowing-*that*, in terms of practically *implicit* capacities, abilities, or knowing-*how*. This would discharge one of the primary explanatory obligations of the pragmatist foe of the intellectualist understanding of norms. For if practical knowing-how is taken as prior in the order of explanation to theoretical knowing-that, one must not only offer an independent account of the practically implicit grasp or mastery of norms. One must also explain how the propositionally explicit grasp of norms expressed in the form of rules, principles, or claims can be understood as arising out of those practical capacities.

In the same way, the inferentialist approach to content treats material proprieties of inferences as prior in the order of explanation to formal logical proprieties of inference. It is accordingly obliged not only to offer an independent account of those material proprieties but also to offer an account of how logical goodness of inference can be explained in terms of that primitive sort of goodness of inference. One who denies that logic is to be understood as underlying (and so presupposed by) rationality in the sense involved in the inferential articulation of conceptual contents (and so in any exercise of the capacity to give and ask for reasons) is obliged to offer another account of logic. This obligation is discharged by the combination of two moves. The first is offering a criterion of demarcation for logical vocabulary that is couched in terms of the semantically expressive role played by such vocabulary in making implicitly content-conferring inferential commitments explicit in the form of judgments. This move depends on having a view about what it is for something to be explicit in the form of a judgeable, that is propositional, content. Such a view is precisely what the account of propositional contents in terms of material proprieties implicit in inferential practice, mentioned above, is intended to supply. The second element required to discharge the obligation to show how the notion of logically good inferences grows out of that of materially good inferences is the substitutional account of formal logical validity of inference—according to which an inference is valid or good in virtue of its logical form if it is primitively good and cannot be turned into one that is not primitively good by any (grammatical) substitution of nonlogical for nonlogical vocabulary.

An explanatory demand exhibiting the same structure as that just rehearsed for the anti-intellectualist about norms and the antiformalist about logic is incumbent on the inferentialist account of conceptual content in virtue of its commitment to invert the representationalist order of semantic explanation.<sup>90</sup> A viable working-out of the inferentialist order of explanation

must, to begin with, offer an account of correctness of inference that is not parasitic on correctness of representation. This demand is addressed in Chapters 3 and 4, which specify sufficient conditions for an attribution of implicitly normative social practices to a community to count as interpreting them as engaging in practices of giving and asking for reasons—as practically assessing inferences as correct or incorrect, and so as instituting material inferential proprieties that confer propositional conceptual content on their states and performances. It is not enough, however, for the inferentialist explanatory strategy to produce an account of the pragmatic basis of its own semantic primitives that does not rely on the prior intelligibility of representational concepts. It must also show how representational relations and the sorts of representational content they underwrite can be made intelligible in terms of those inferential primitives. That is, another critical criterion of adequacy of inferentialism is the extent to which, if this approach is granted its preferred starting point, it can develop it into an account of the sort of objective representational content other approaches begin with.

### 3. *Objective Representational Content*

Meeting this demand involves offering accounts of three important dimensions along which the notion of objective representational content is articulated. First is the referential dimension. The representationalist tradition has, beginning with Frege, developed rich accounts of *inference* in terms of *reference*. How is it possible conversely to make sense of reference in terms of inference? In the absence of such an account, the inferentialist's attempt to turn the explanatory tables on the representationalist tradition must be deemed desperate and unsuccessful.

The second dimension is categorial. An account must be offered not just of reference and representation but of reference to and representation of particular *objects* and general *properties*. That is, the peculiar kind of representational content expressed by *subsential* expressions, paradigmatically singular terms and predicates, must be explained. For reasons already indicated, inferential approaches to conceptual content apply directly only to what is expressed by declarative sentences, which can play the role of premises and conclusions of inferences. Somehow the inferential approach to conceptual content must be extended to apply to subsential parts of speech as well. The discussion of Dummett's model of circumstances and consequences of application provides some suggestive hints. But these must be developed far beyond the remarks already offered in order to put the inferentialist in a position to claim to have shown that the nominalist order of explanation standard prior to Kant, beginning with a doctrine of terms or concepts and moving from there to a doctrine of judgments, can successfully be stood on its head.

Finally, there is the *objective* dimension of representational content. It



must be shown how on inferentialist grounds it is possible to fund *objective* proprieties of inferring and judging—to make intelligible the way in which what it is correct to conclude or to say depends on how the objects referred to, talked about, or represented actually are. Even if, to begin with, attention is restricted to inferential proprieties, it is clear that not just any notion of correctness of inference will do as a rendering of the sort of content we take our claims and beliefs to have. A semantically adequate notion of correct inference must generate an acceptable notion of conceptual content. But such a notion must fund the idea of *objective* truth conditions and so of *objectively* correct inferences. Such proprieties of judgment and inference outrun actual attitudes of taking or treating judgments and inferences as correct. They are determined by how things actually are, independently of how they are taken to be. Our cognitive attitudes must ultimately answer to these attitude-transcendent facts.

This means that although the inferentialist order of explanation may start with inferences that are correct in the sense that they are accepted in the practice of a community, it cannot end there. It must somehow move beyond this sense of correctness if it is to reach a notion of propositional conceptual content recognizable as that expressed by our ordinary empirical claims and possessed by our ordinary empirical beliefs. Pursuing the inferentialist order of explanation as outlined above accordingly requires explaining how—if actual practical attitudes of taking or treating as correct institute the normative statuses of materially correct inferences, and these material proprieties of inference in turn confer conceptual content—that content nonetheless involves objective proprieties to which the practical attitudes underlying the meanings themselves answer.<sup>91</sup> How is it possible for our use of an expression to confer on it a content that settles that we might all be wrong about how it is correctly used, at least in some cases? How can normative attitudes of taking or treating applications of concepts as correct or incorrect institute normative statuses that transcend those attitudes in the sense that the instituting attitudes can be assessed according to those instituted norms and found wanting? This issue of objectivity is perhaps the most serious conceptual challenge facing any attempt to ground the proprieties governing concept use in social practice—and the pragmatist version of inferentialism being pursued here is a view of this stripe.

In the terms set up in Section I of this chapter, the referential, categorial, and objective can be thought of as three interlocking dimensions of the project of explaining *object-representing* contentfulness in terms of *propositional* contentfulness, according to a semantic rendering of propositional contentfulness in terms of *material proprieties of inference* and a *pragmatic* rendering of those basic inferential proprieties.<sup>92</sup> The relation between inference and reference is discussed in an introductory way in Chapter 5, which examines the use of the semantic vocabulary (paradigmatically 'refers' and 'true') by whose means the implicit referential dimension of conceptual

contents is made explicit. The anaphoric relations that are invoked in the explanation offered there are then explained in more primitive pragmatic terms in Chapter 7, which relates them to the sort of substitution-inferential commitments discussed in Chapter 6. The categorial issue is addressed by further development of Frege's substitutional methodology, in Chapter 6, which shows how the contents expressed by the use of singular terms and predicates can be understood in terms of substitution inferences.

The objectivity issue, which concerns the relation between what is properly *said* and what is talked *about*, intimately involves both the referential or representational dimension and the categorial. As with all substantive semantic and pragmatic concepts officially employed in this work in describing the contents and significances of discursive commitments, the concept of states, attitudes, and performances that are objectively contentful in purporting to represent how things are independently of anyone's states, attitudes, and performances is discussed at two levels. (This two-leveled account is part of the effort to secure for the use of such vocabulary in this project an analog of the kind of *expressive equilibrium* already adverted to as achieved by Frege's treatment of sentential logical vocabulary in the *Begriffsschrift*.) The first is a story about what it is for such purport and its uptake to be *implicit* in the practices of those whose states, attitudes, and performances are properly interpreted as having such content.

The second is a story about what it is for such purport and its uptake to be made *explicit* in the specification of contents of ascribed states, attitudes, and performances. Although technical philosophical vocabulary such as 'refers' or 'denotes' (discussed in Chapter 5) can play this explicating role, the fundamental locutions used in ordinary talk to express representational commitments are those used to form *de re* specifications of the contents of ascribed intentional states, attitudes, and performances—paradigmatically 'of' and 'about'. The use of *de re* ascriptions makes it possible to specify explicitly what is said in terms of what is talked about. What such ascriptions express and how those objective content-specifications are made explicit by their use is the topic of Chapter 8. The account of what we are doing when we interpret ourselves and each other as making claims with objective representational content that is offered in that chapter requires the expressive resources of all of the sorts of locutions whose use is introduced in prior chapters. It is accordingly only in the last substantive chapter of this work that this critical explanatory obligation of an inferential approach to semantics is finally discharged.

The foundation of that account is laid in the next chapter. It consists in the *social* structure of the inferential norms that confer propositional content. (Government by such norms is what such contentfulness is.) The development into a more full-blooded notion of conceptual content of the abstract notion of inferential role introduced in this chapter proceeds by taking account of the social dimension of inferential practice—which is

implicit in the idea that abstract talk about inferential relations must be rooted in consideration of what Sellars calls "the game of giving and asking for reasons." The pragmatic significance of making a claim or acquiring a commitment whose content could be expressed by the use of a declarative sentence cannot be determined by associating with that sentence a set of sentences that entail it and a set of sentences that it entails—not even if these are enriched by throwing in nonlinguistic circumstances and consequences of application as well. This is because of the interaction of two features of inferentially articulated commitments.

First, as Frege acknowledges in his original definition of *begriffliche Inhalt*, specification of the inferential role of a sentence requires looking at *multipremise* inferences.<sup>93</sup> Many of the important "consequences of application" of a sentence are not consequences it has all on its own; they consist rather in the differential contribution its inclusion makes to the consequences of a set of collateral premises or auxiliary hypotheses. Similarly, its purely inferential antecedents must be thought of not as individual sentences but as sets of them.

Second, the collateral concomitant commitments available as auxiliary hypotheses in multipremise inferences vary from individual to individual (and from occasion to occasion or context to context). If they did not, not only the notion of communication but even that of empirical information would find no application. The significance of acquiring a commitment or making a claim whose content could be expressed by the use of a particular sentence, when it would be appropriate to do so and what the appropriate consequences of doing so would be, depends on what other commitments are available as further premises in assessing grounds and consequences. What is an appropriate ground or consequence of that commitment from the point of view of one set of background beliefs may not be from the point of view of another. In view of the difference in their other attitudes, a single commitment typically has a different significance for the one undertaking it, a speaker or believer, from that which it would have for those attributing it, an audience or intentional interpreter. Of course this does not make communication or interpretation impossible—on the contrary. As was just mentioned, it is only the prevalence of situations in which background commitments do differ that give communication and interpretation their point.

The fact that the implicitly normative inferential significance of a commitment may be different from the point of view of one undertaking the commitment and one attributing it means that the inferential articulation of conceptual contents has a fundamental *social* dimension. It introduces a relativity to social perspective into the specification of such contents. The practical attitudes of taking or treating as committed, which ultimately institute the normative status of commitment, come in the two socially distinct flavors of undertaking or acknowledging a commitment (oneself) and

attributing a commitment (to another). Inferentially articulated contents are conferred on states, attitudes, and performances by the norms instituted by *social* practices: those that essentially involve the interaction of attitudes corresponding to both social perspectives. Investigation of the use of locutions that make explicit various aspects of the social perspectival character of conceptual contents will reveal what they express as the source of objective representational content. So, it will be claimed, what must be added to the *normative* approach to pragmatics and the *inferential* approach to semantics in order to make intelligible the *representational* dimension of conceptual contents is a *social* account of the interaction between them.

## Linguistic Practice and Discursive Commitment

---

Language most shows a man: Speak, that I may see thee.

BEN JONSON, *Timber or Discoveries*

Language is called the Garment of Thought: however, it should rather be, Language is the Flesh-Garment, the Body, of thought.

THOMAS CARLYLE, *Sartor Resartus*

Clearly human beings could dispense with all discourse, though only at the expense of having nothing to say.

WILFRID SELLARS, "A Semantical Solution to the Mind-Body Problem"

---

### I. INTENTIONAL STATES AND LINGUISTIC PRACTICES

#### 1. *Discursive Practice as Deontic Scorekeeping*

This chapter introduces a particular model of language use: the deontic scorekeeping model of discursive practice. The implicitly normative social practices it describes are inferentially articulated in such a way as to confer specifically *propositional* contents on expressions and performances that play suitable roles in those practices. The basic idea is the one motivated by the discussion in Chapter 2, namely that propositional contentfulness must be understood in terms of practices of giving and asking for *reasons*. A central contention is that such practices must be understood as *social practices*—indeed, as *linguistic practices*. The fundamental sort of move in the game of giving and asking for reasons is making a *claim*—producing a performance that is propositionally contentful in that it can be the offering of a reason, and reasons can be demanded for it. Other theoretically important concepts are defined in terms of this one: *linguistic practice* is distinguished by its according some performances the significance of claimings, and (declarative) *sentences* are distinguished as expressions whose utterances, inscriptions, or other tokenings have the default significance of claimings. The basic explanatory challenge faced by the model is to say what structure a set of social practices must exhibit in order properly to be understood as includ-

ing practical attitudes of taking or treating performances as having the significance of claims or assertions.

According to the model, to treat a performance as an assertion is to treat it as the undertaking or acknowledging of a certain kind of *commitment*—what will be called a ‘doxastic’, or ‘assertional’, commitment. To be doxastically committed is to have a certain social status. Doxastic commitments are normative, more specifically *deontic*, statuses. Such statuses are creatures of the practical attitudes of the members of a linguistic community—they are instituted by practices governing the taking and treating of individuals *as* committed. Doxastic commitments are essentially a kind of deontic status for which the question of *entitlement* can arise. Their inferential articulation, in virtue of which they deserve to be understood as propositionally contentful, consists in consequential relations among the particular doxastic commitments and entitlements—the ways in which one claim can commit or entitle one to others (for which it accordingly can serve as a reason).

Competent linguistic practitioners keep track of their own and each other’s commitments and entitlements. They are (we are) *deontic scorekeepers*. Speech acts, paradigmatically assertions, alter the deontic score; they change what commitments and entitlements it is appropriate to attribute, not only to the one producing the speech act, but also to those to whom it is addressed. The job of pragmatic theory is to explain the *significance* of various sorts of speech acts in terms of practical proprieties governing the keeping of deontic score—what moves are appropriate given a certain score, and what difference those moves make to that score. The job of semantic theory is to develop a notion of the *contents* of discursive commitments (and the performances that express them) that combines with the account of the significance of different kinds of speech act to determine a scorekeeping kinematics.

The basic elements of this deontic scorekeeping model of discursive practice are presented in this chapter. The next chapter develops it further by attending in particular to the inferential articulation of *perception* and *action*. These are the entries to and exits from the realm of discursive commitments and entitlements—the source respectively of the empirical and practical dimensions of conceptual content, which are usually (and in one sense correctly) thought of as *noninferential*. Before plunging into a description of the details of the features of a system of social practices in virtue of which they should be understood as incorporating practical scorekeeping attitudes that institute deontic statuses and confer propositional contents, however, some methodological preliminaries are in order. The rest of this work presents not only an account of linguistic intentionality (thought of as one sophisticated species in a genus comprising other, more primitive sorts) but a linguistic account of intentionality generally. It is claimed that the propositional contentfulness even of the beliefs and other states intentional

interpreters attribute to nonlinguistic animals cannot properly be understood without reference to the specifically linguistic practice of the interpreters, from which it is derived. Original, independent, or nonderivative intentionality is an exclusively linguistic affair. The reasons for insisting on the conceptual primacy of linguistic intentionality cannot be presented until all the materials needed for the analysis of the representational dimension of propositional content (and of conceptual content generally) have been assembled, in Chapter 8. Nonetheless, the explanatory strategy being pursued will be easier to understand if the picture of the relations between language and belief that it incorporates has been sketched, even if the warrant for that picture cannot emerge until it is more fully developed (in Part 2 of this work).

## 2. *Philosophical Semantics and Formal Semantics*

One of the fundamental methodological commitments governing the account presented here is *pragmatism* about the relations between semantics and pragmatics. Pragmatism in this sense is the view that what attributions of semantic contentfulness are *for* is explaining the normative significance of intentional states such as beliefs and of speech acts such as assertions. Thus the criteria of adequacy to which semantic theory's concept of content must answer are to be set by the pragmatic theory, which deals with contentful intentional states and the sentences used to express them in speech acts. The idea that *philosophical* theories of meaning or content must be concerned with the larger pragmatic context within which attributions of contentfulness play an explanatory role may seem to be brought into question by the autonomy of *formal* semantics. But the independence of formal semantics from pragmatic concerns is only apparent.

The project of formal semantics entitles the theorist at the outset to stipulate an association of semantic interpretants with primitive interpreteds, typically linguistic expressions. Then this interpretation is extended to interpreteds that are derived from those primitives by syntactic operations—which for standard compositional syntactic structures include category-sensitive concatenation and various grammatical transformations of such concatenations. This is achieved by defining, for each syntactic operation on interpreteds, a corresponding operation on their associated interpretants that yields a new interpretant, which is thereby associated with the result of the syntactic operation. So the formal semantic theorist might begin by associating truth-values with sentence-letters, and then for each connective that produces compound sentences introduce a function taking sets of truth-values into truth-values that can then be assigned to the corresponding compound sentences. Or instead of truth-values, the semantic interpretants might be sets of possible worlds, and the operations corresponding to sentential connectives be set-theoretic operations on them (such as intersection for conjunction).<sup>1</sup>

So formal semantics is concerned generically with structure-preserving mappings. But not every mathematical representation theorem, which shows such a correspondence between structures of one kind and those of another, deserves to be called a semantics. What else ought to be required for a set of such mappings to count as presenting a specifically *semantic* interpretation of something? To ask this question is already to begin to move from the domain of purely formal semantics to that of philosophical semantics. When Tarski proved an algebraic representation theorem in which the interpretants assigned to quantificational expressions are topological closure operators, what qualifies that as a formal semantics for the first-order predicate calculus is not anything about the intrinsic properties of those interpretants but just that he is able in those terms to reproduce the relation of logical consequence appropriate to that idiom. From a purely formal or mathematical point of view, the task would be no different if the property to be reproduced were specified simply by randomly partitioning the elements of one grammatical category, placing an asterisk next to some of them and not others (and similarly for the relation in question). From the point of view of the philosophical motivation of calling what one is doing thereby 'semantics', however, it makes all the difference that the elements involved be interpretable as *sentences*, and that the property distinguished be interpretable as theoremhood, a kind of *truth*, and that the relation distinguished be interpretable as derivation, a kind of *inference*. Indeed, to take the elements as subject to evaluations concerning propriety of *judging* and propriety of *reasoning*, truth and inference, is just what it is to interpret them as sentences.

What gives semantic theory its philosophical point is the contribution that its investigation of the nature of contentfulness can make to the understanding of proprieties of practice, paradigmatically of judging and inferring. That semantic theory is embedded in this way in a larger explanatory matrix is accordingly important for how it is appropriate to conceive the semantic interpretants associated with what is interpreted. It means that it is pointless to attribute semantic structure or content that does no pragmatic explanatory work. It is only insofar as it is appealed to in explaining the circumstances under which judgments and inferences are properly made and the proper consequences of doing so that something associated by the theorist with interpreted states or expressions qualifies as a *semantic* interpretant, or deserves to be called a theoretical concept of *content*. Dummett puts the point this way:

The term 'semantics', at least as commonly applied to formalized languages, usually denotes a systematic account of the truth-conditions of sentences of the language: the purpose of thus assigning a value, true or false, to every well-formed sentence of the language is taken as already understood, and receives no explanation within the semantic theory itself . . . The classification of the sentences of a formalized



language into true ones and false ones relates to the purposes for which we want to use the language. But in the case of natural language, it is already in use: the only point of constructing a semantics for the language can be as an instrument for the systematic description of that use, that is, as part of a whole theory of meaning for the language, which as a whole constitutes an account of its working. If the semantic part of the theory is taken as issuing an assignment of conditions under which each sentence of the language, as uttered on a particular occasion, has this or that truth-value, the rest of the theory must connect the truth conditions of the sentences with the use to which they are put, that is, with the actual practice of speakers of the language . . . a semantic theory which determines the truth-conditions of sentences of a language gets its point from a systematic connection between the notions of truth and falsity and the practice of using those sentences.<sup>2</sup>

The essential point is that philosophical semantic theory incorporates an obligation to make the semantic notions it appeals to intelligible in terms of their pragmatic significance. Formal semantics qualifies as *semantics* only insofar as it is implicitly presupposed that this obligation can be satisfied by conjoining the semantics with some suitable pragmatics.<sup>3</sup>

Philosophical semantics is distinguished from formal semantics by its explicit concern with the relation between the use of semantic concepts, on the one hand, and pragmatic accounts of the proprieties of practice governing the employment of what those concepts apply to, on the other. Philosophical semantics is committed to explaining the content of concepts such as *content*, *truth*, *inference*, *reference*, and *representation*, while formal semantics is content to use such concepts, assuming them (and so the pragmatic significance of applying them) already to be implicitly intelligible. The difference between doing either sort of semantics for artificial languages and for natural languages is that in the former case there are no antecedent proprieties governing the use of the expressions, to which the semantic theorist is responsible. Since the language is not already in use, the theorist is free to *stipulate* an association of contents with expressions, in order to determine how they are to be understood to be correctly used. In the case of natural languages, however, the theorist's use of semantic concepts is not *synthetic* (to settle the proper employment of expressions that antecedently are subject to no such proprieties) but *analytic* (to codify and express antecedently existing proprieties of employment).

### 3. *Associating Content Explicitly by Stipulation and Implicitly Conferring It by Practice*

Philosophical semantic theories of expressions and states that already play normatively articulated roles in linguistic practice or in the

practical reasoning of rational agents accordingly cannot afford the luxury (enjoyed by formal semantics of all sorts and even by philosophical semantics of artificial languages) of employing a *stipulative* method. Such theories are obliged to explain what the association of content with expression or state *consists in*: what one is saying or doing in attributing content to them. At this point it has seemed to many that the cases of contentful *sentences* and of contentful *beliefs* diverge. It makes sense to think of the contents of linguistic expressions as *conferred* on them by the way they are used. Noises and marks mean nothing all on their own. No one thinks they are *intrinsically* contentful. The sign-design 'dog' could as well be used to express the concept expressed by the sign-design 'horse', or to express none at all, like 'gleeb'. It is only by being caught up in linguistic practice that they come to express propositions, make claims, have or express conceptual or intentional contents. Apart from their role in human activity, apart from the norms thereby imposed on their employment—which make it the case that some uses are correct and others incorrect—these linguistic vehicles are semantically mute, inert, dead.

The philosophical semantics of natural languages must begin, then, with the observation that it is the practice of those who use the language that confers content on the utterances and inscriptions that are the overt, explicitly expressive performances whose propriety is governed by that practice. Is something similar true of intentional states? There are some important asymmetries between the two cases. There is a familiar line of thought, already adverted to, according to which quite a different story must be told about the association of content with the states and attitudes, paradigmatically beliefs, that are expressed by such linguistic performances. The critical question is how to understand the *use* of language in which the pragmatic significance of speech acts consists and which accordingly confers semantic content on those speech acts and so indirectly on the expressions they involve. One way to think about such use is *instrumentally*. This line of thought may be traced back to Locke, who thought of speech as an instrument for communicating thoughts or ideas. It is successful when the noise emitted by the speaker arouses in the audience an idea with the same content as that prompting the speech act.

Contemporary elaborations of this approach see "nonnatural" meaning as rooted in the capacity of individuals deliberately to imbue signals with significance by producing them with the intention that they be understood in a certain way by their auditors. According to Grice's picture,<sup>4</sup> linguistic practitioners *make* their expressions have a certain content by producing them with the intention that others *take* them to have that content. In particular, assertion is understood as the expression of belief in the sense that a sentence is produced with the intention that those who hear it will acquire a certain belief in virtue of their recognition that the speaker uttered the sentence intending those who hear it to acquire that belief, in virtue of their

recognition of that very intention. The notion of a linguistic expression's meaning something is in this way derived from the capacity of language users deliberately to mean something by their utterances. In somewhat different ways, Lewis, Bennett, and Searle develop this instrumental model by showing how a layer of *conventions* can be built on such communicative *intentions*, in such a way that members of a linguistic community are for the most part relieved of the necessity for elaborate deliberation about each other's beliefs and intentions in choosing and interpreting each other's remarks.<sup>5</sup> The foundation on which the conventional meaningfulness of linguistic acts and expressions rests remains their intentional employment as means to an explicitly envisaged communicational end.<sup>6</sup>

Rosenberg calls this explanatory strategy "agent semantics," because linguistic meaning is explained in terms of a prior capacity to engage in practical reasoning.<sup>7</sup> If the pragmatic use of language that confers semantic content on utterances and expressions is understood in these terms, it is clear that the contentfulness of intentional states such as beliefs, desires, and intentions must be understood antecedently, and hence according to some other model. Agent semantics treats the contentfulness of utterances as derivative from that of intentional states. The content of an assertion derives from the content of the belief it is the expression of, and from the content of the intention that it be understood as expressing that belief. It follows that it must be possible to make sense of the contents of beliefs and intentions prior to and independently of telling this sort of story about the use of linguistic expressions. Their content cannot be taken to be conferred on them by the way they are used or employed, according to this model of use or employment.<sup>8</sup>

That the content of intentional states cannot be understood as conferred on them by proprieties governing their significance—when it is appropriate to acquire them and what the appropriate consequences of acquiring them are—follows only if the only candidate for content-conferring use is deliberate, instrumental employment in order to secure the *explicitly* envisaged purpose of being understood as having a certain content. It would not follow that semantic content could not be conferred on intentional states by proprieties *implicit* in the way those states are treated in practice. According to such a conception, the conferral of content might be a side effect of the way they are treated, not requiring that anyone explicitly intend to confer it by their behavior. Broadly functionalist approaches to content are of this sort. They understand intentional states to be contentful in virtue of the role they play in the proper functioning of some system of which they are a part. Going into a certain state is something that is done appropriately under some circumstances, according to the functional interpretation of the system, and it has certain appropriate consequences. Together these proprieties of input and output, antecedents and consequences, determine the functional role of the state in the system. According to the functionalist explanatory strategy,

it is in virtue of playing this role, being subject to these transitional proprieties, that intentional states have the content they do.

#### 4. *Intentionality: Linguistic Practice versus Rational Agency*

The first question that needs to be addressed in working out such an approach is how the relevant functional system should be understood. Is it possible to understand propositional and other genuinely conceptual contents as conferred on states and performances by their role in a functional system that comprises only a single individual? Or is discursive practice essentially *social* practice, so that the functional system must be taken to comprise the activities of an entire community? The most popular and promising way of developing the first answer looks to the role belief plays in the practical instrumental reasoning of intelligent agents. The most popular and promising way of developing the second answer looks to the role assertion—the explicit expression of belief—plays in linguistic practice.

The considerations assembled in the first two chapters suggest the motivation that these two approaches have in common: states, attitudes, and performances are intentionally contentful in virtue of the role they play in *inferentially* articulated, implicitly *normative* practices. It is by looking at the practices in which the status of some states, attitudes, and performances as providing *reasons* for others is implicitly (and constitutively) acknowledged that the pragmatic significance of associating them with intentional contents is to be understood. There are two different sorts of context in which the specifically inferential significance of intentional states such as belief is to the fore: rational agency and linguistic practice. On the one hand, beliefs and other intentional states are expressed in *actions*, nonlinguistic performances that are intelligible in virtue of the beliefs and desires that are reasons for them. On the other hand, beliefs are expressed in *claims*. Overt assertions are the fundamental counters in the game of giving and asking for reasons—they can both be offered as reasons and themselves stand in need of such reasons.

Stalnaker points out that these two sorts of context in which intentional states are significant for practice generate two basically different ways of looking at those states: what he calls the *pragmatic* picture of intentionality and the *linguistic* picture of intentionality. The pragmatic picture is one according to which “rational creatures are essentially agents . . . According to this picture, our conceptions of belief and of attitudes pro and con are conceptions of states which explain why a rational agent does what he does . . . Linguistic action, according to this picture, has no special status. Speech is just one kind of action which is to be explained and evaluated according to the same pattern. Linguistic action may be a particularly rich source of evidence about the speaker’s attitudes, but it has no special conceptual connection with them.”<sup>9</sup> This picture amounts to a generalization of the ap-

proach of agent semantics. It shares the emphasis on rational agency as the conceptually and explanatorily fundamental context in which to understand the significance of the contentfulness of intentional states. It is more general in that it does not necessarily involve a commitment to understanding the contentfulness of speech acts as deriving from their deliberate instrumental employment to secure antecedently envisageable goals. Thus this picture leaves room for a picture in which the way speech acts inherit the contents of the intentional states they express might be less intellectualized, the conferral of content being implicit in the practice of expression, rather than explicit as its instrumentally conceived motive. The contrasting linguistic picture is one according to which "rational creatures are essentially speakers . . . Representational mental states represent the world because of their resemblance to or relation with, the most basic kind of representations: linguistic expressions."<sup>10</sup>

This way of dividing up the fundamental orientations of various approaches to intentionality, accordingly as rational agency or linguistic capacity is taken as primary, evidently cuts at important joints. It is a measure of the robustness of this botanization that it is serviceable even across large differences in collateral theoretical commitments. Here is how Stalnaker puts the setting in which he sees the pictures as competing: "The linguistic and pragmatic pictures each suggest strategies for giving a naturalistic explanation of representation—both mental and linguistic representation—but the two strategies differ in what kind of representation they take to be more fundamental. The pragmatic picture suggests that we explain the intentionality of language in terms of the intentionality of mental states, while the linguistic strategy suggests that we explain the intentionality of mental states in terms of, or by analogy with, the intentionality of linguistic expressions."<sup>11</sup> On the side of semantic content, Stalnaker follows the tradition in seeing the issue as fundamentally one of representation, with inferential relations presumably to be explained further along in ultimately representational terms. And on the side of the pragmatic significance of intentional contentfulness, Stalnaker begins with a commitment to a naturalistic approach, with the normative character of the practice in which intentional states are significant (whether that practice is conceived in the first instance as rational agency or as essentially linguistic) presumably to be explained further along in ultimately naturalistic terms.

The order of explanation that frames his discussion is the reverse of that pursued in this work. For the semantic explanatory strategy being developed here looks first to inference, on the semantic side, and aspires to making the representational dimension of intentional content intelligible ultimately in inferential terms. And on the pragmatic side, the strategy is to begin with an account of norms implicit in practice and work out toward an understanding of their relation to their naturalistic setting, which the normative practices in their most sophisticated form make it possible to describe objectively. It

is noteworthy that in spite of these major differences in approach, the large division of options for explaining intentionality into those focusing on rational agency and those focusing on language seems compelling from both points of view.

### *5. Analogical and Relational Versions of Linguistic Approaches to Intentionality*

Broadly linguistic approaches comprise many importantly different variants, however, and these correspond to importantly different motivations. Stalnaker implicitly acknowledges one significant subdivision within the linguistic approach in his general characterization of the linguistic picture as seeking to explain the contentfulness of intentional states by appealing to their “resemblance to or relation with” contentful linguistic performances. The disjunction links two very different ways in which it might be thought that taking account of specifically linguistic practice is essential to making the contentfulness of intentional states explicit (in the sense of theoretically intelligible). The resemblance limb, according to which the contentfulness of states is modeled on that of expressions, involves a commitment to the claim that the theorist’s or interpreter’s capacity to attribute (and understand attributions of) contentful intentional states is parasitic on the theorist’s or interpreter’s capacity to attribute (and understand attributions of) contentful speech acts. It need not entail, for instance, that the intentional states attributed to nonlinguistic creatures are somehow second class. It requires only that what one is doing in attributing contentful states to nonlinguistic creatures cannot be understood apart from the capacity to attribute them to linguistic ones. By contrast, the relational limb—according to which the contentfulness of intentional states consists in or essentially involves the contentfulness of the speech acts that express them—involves a commitment to the claim that the theorist’s or interpreter’s capacity to attribute (and understand attributions of) contentful intentional states is in the first instance parasitic on the theorist’s or interpreter’s capacity to attribute (and understand attributions of) contentful speech acts to the *same* individuals who are taken to have the intentional states. It does entail that the intentional states attributed to nonlinguistic creatures are in important regards second-class statuses.<sup>12</sup>

One reason it is important to distinguish the claim that the intentionality of speech is conceptually prior to the intentionality of belief (as analogical theories have it) from the claim that as believers in the full sense, we are essentially rather than only accidentally speakers (as relational theories have it) is that only theories committed to the former thesis are obliged to offer accounts of linguistic practice that do not make reference to intentional states. It is open to one who subscribes to the second view to hold, as Davidson does, that attributions of contentful intentional states and content-

ful speech acts go hand in hand, that neither sort of attribution is intelligible apart from its relation to the other. So a relational account can understand the possibility of speech as essential to intentionality (in the paradigm cases from which we derive our grip on what intentional interpretation is) without thereby becoming obliged simply to invert the order of explanation characteristic of agent semantics—though such an account evidently cannot appeal to the sort of independently intelligible role of contentful states in rational agency presupposed by agent semantics.

In contrast, analogical linguistic theories of intentionality are committed to that converse order of explanation. Agent semantics employs an antecedent and independent notion of the contentfulness of intentional states to explain the derivative contentfulness of speech acts and linguistic expressions. A theory insisting that the contentfulness of intentional states is intelligible only by analogy to the contentfulness of speech acts and linguistic expressions would be obliged correspondingly to appeal to an antecedent and independent notion of the contentfulness of speech in order to explain the derivative contentfulness of intentional states. A relational linguistic theory of intentionality maintains rather that the understanding of intentionally contentful states that permits us to stretch the application of that notion and apply it in a second-class way to nonlinguistic animals (simple intentional systems) derives from and essentially depends on an understanding of the relation between the intentional states and the linguistic performances of language-using animals (communicating or interpreting intentional systems).

According to this sort of approach, understanding intentionality requires looking at practices that essentially involve *both* intentional states and linguistic performances. Neither sort of intentionality need be understood as conceptually prior to the other, and linguistic practice and rational agency can be presented as two aspects of one complex of jointly content-conferring practices. Davidson puts the characteristic contention of relational linguistic views of intentionality this way: "Neither language nor thinking can be fully explained in terms of the other, and neither has conceptual priority. The two are, indeed, linked, in the sense that each requires the other in order to be understood; but the linkage is not so complete that either suffices, even when reasonably reinforced, to explicate the other."<sup>13</sup> The account of intentionality introduced in the rest of this chapter is a linguistic theory in this relational sense. The key to motivating a theory of this sort is to show what it is about the contents of intentional states that can be explained only by appealing to the relation between such states and specifically *linguistic* performances.

Davidson suggests that an argument for a relational theory can be provided in two pieces. He claims first that "someone cannot have a belief unless he understands the possibility of being mistaken, and this requires grasping the contrast between truth and error—true belief and false belief."<sup>14</sup> What a

creature has does not function as a belief for that creature unless it has a certain kind of significance for that creature. It must be able to adopt a certain kind of practical attitude toward that state, to treat it in its practice or behavior *as* contentful in a special sense. In particular, Davidson is claiming, what it has is not recognizable as belief unless the creature whose state it is somehow in its practice acknowledges the applicability of a distinction between beliefs that are correct and those that are incorrect, in the sense of being true and false. We are not permitted to attribute the belief that *p* (a propositionally contentful intentional state) unless somehow the putative believer acknowledges in practice the *objective representational* dimension of its content—that its being held is one thing, but its being *correct* is another, something to be settled by how it is with what it is *about*. The second piece of Davidson's argument is the claim that a grasp of the contrast between correct and incorrect belief, true and false belief, "can emerge only in the context of interpretation, which alone forces us to the idea of an objective, public truth." The key claim is that "the concepts of objective truth and of error necessarily emerge in the context of interpretation."

The rest of this work focuses on the development of an account of linguistic social practices within which states, attitudes, and performances have, and are acknowledged by the practitioners to have, pragmatic significances sufficient to confer on them objective representational propositional contents. The view propounded is like Davidson's in seeing intentional states and speech acts as fundamentally of coeval conceptual status, neither being explicable except in an account that includes the other. It deserves nonetheless to be called a *linguistic* view of intentionality (of the relational rather than the analogical variety) because linguistic practice is nonetheless accorded a certain kind of explanatory priority over rational agency. The intentionality of nonlinguistic creatures is presented as dependent on, and in a specific sense derivative from, that of their linguistically qualified interpreters, who as a community exhibit a nonderivative, original intentionality. The sort of derivation in question is explicated in terms of the context that must be appealed to in making intelligible the sort of contents (*conceptual*—paradigmatically *propositional*—contents) that are associated with the intentional states attributed by interpreters. The contents of the intentional states attributed to nonlinguistic creatures can be understood only in a way that involves the activities of the language users who attribute them, and not entirely in terms of the activities of those who exhibit them. By contrast, the contents of the intentional states attributed to a community of language users can be understood as conferred on their states, attitudes, and performances entirely by the practices of those community members.

The argument that provides the ultimate justification for treating specifically *linguistic* practice as central in this way to intentionality has just the two-part form outlined by Davidson and rehearsed above. For what he



has really given us is not so much an argument as the form of one. Turning it into an actual argument requires filling in various notions of content, of objective representational correctness of content, of practical acknowledgment of the significance of assessments of correctness of content, and so on. That is the task of the rest of this work; the final justification for giving pride of place to language will thus not be complete until the end (in fact, in Chapter 8). At that point it will be possible to return again to the beginning, and know it for the first time—the warrant for this fundamental theoretical commitment will then be explicit.

### 6. *Believing and Claiming*

Claiming and believing are linked by the principle that assertions are one way of expressing beliefs. A fundamental question of explanatory strategy is then whether this principle can be exploited so as to account for one of these notions in terms of the other. Since there can be beliefs that are not avowed, the temptation is to start with belief and explain assertion as a speech act by which belief is expressed. But when the representational dimension of propositional and other conceptual contents is examined (in Chapter 8), it turns out to be intelligible only in the context of *linguistic* social practices of communicating by giving and asking for reasons in the form of claims. So if assertion were to be explained in terms of a prior notion of belief, the propositional contents of beliefs would have to be taken for granted. Their association with beliefs would have to be stipulated, rather than made intelligible as established by the functional role of beliefs in the behavioral economies of believers. The only sort of inferential practice that is socially articulated in the way that turns out to be required for the conferral of *propositional* content, in the form of objective truth conditions, is *assertional*, and therefore *linguistic* practice.

The idea pursued here is that the state or status of *believing* is essentially, and not merely accidentally, related to the linguistic performance of *claiming*. Beliefs are essentially the sort of thing that can be expressed by making an assertion. Dummett offers a crisp formulation of a view along these lines: "We have opposed throughout the view of assertion as the expression of an interior act of judgment; judgment, rather, is the interiorization of the external act of assertion."<sup>15</sup> Although the satisfyingly symmetric phrasing of this remark can obscure the point, it is important to realize that Dummett is not committed by it to the possibility of making sense of the speech act of asserting without mentioning anything but speech acts. For instance, this stance does not preclude an account of asserting that incorporates an account of the particular sort of *commitment* (a deontic status) one undertakes in making an assertion. What is precluded is only explaining assertion as the expression of a kind of intentional state or deontic status that is supposed to

be intelligible apart from the possibility of expressing it by asserting something.

The claim is, then, that speech acts having the pragmatic significance of assertions play an essential role in (social) functional systems within which states or statuses can be understood as propositionally contentful in the way beliefs are. A good way to see how explanatory progress might be made in thinking about beliefs by insisting on their linguistic expression in claims is provided by combining Dummett's thought with a suggestion of Harry Field's concerning how an appeal to language might function as part of a divide-and-conquer strategy for explaining intentional states. His idea is that having a belief with a certain propositional content should be understood as standing in a certain relation to a sentence that expresses that content.<sup>16</sup> In particular, according to what he calls the "two-stage" explanatory strategy, *A* believes that *p* if and only if there is a sentence  $\sigma$  such that:

1. *A* believes\*  $\sigma$ , and
2.  $\sigma$  means that *p*.

In the first stage of an account with this structure, the relevant relation between a believer and a sentence—what Field calls "belief\*"—must be explained. In the second stage, what it is for a sentence to express a propositional content must be explained. Field's own way of pursuing this strategy takes the sentences involved to belong to a "language of thought," which is conceived by analogy to public languages. This is an additional theoretical commitment on his part; nothing about the two-stage decomposition dictates that the sentences that play the role of middle terms should not be sentences in a public language such as English. Filling in the appeal to sentences by invoking a language used in interpersonal communication makes the two-stage strategy available for duty in what were called above relational linguistic theories of intentionality.

In this form, Dummett's claim about the fundamental importance of the speech act of assertion can be pressed into service in addressing the first subproblem of the two-stage strategy. One way in which beliefs are manifested or expressed is by the utterance of sentences. Sentence-utterances can have many sorts of force or pragmatic significance, but when such performances have the significance of *assertions*, they express or purport to express beliefs. As Dummett suggests, this fact can be exploited by two different orders of explanation. If the theorist can get an independent grip on the notion of belief, typically from the consideration of its functional role in the sort of practical reasoning implicitly attributed by intentional interpretations that use the model of rational agency to make nonlinguistic performances instrumentally intelligible, then asserting might be explained in terms of it, as a further way in which beliefs can be manifested. Or, if the theorist can get an independent grip on the practices within which performances are accorded the significance characteristic of assertions, belief might be ex-

plained as what is thereby expressed. If this path is followed, then only parties to linguistic practices, those that institute assertional significances, will qualify as believers. The decision as to which direction of explanation to adopt is in part an issue concerning just how important language is to us.

Like all fundamental demarcational matters, however, it is only in part a factual issue. There clearly is a sense in which nonlinguistic animals can be said to have beliefs. But the sense of belief that Sellars, Dummett, and Davidson are interested in (and that is the subject of this work) is one in which beliefs can be attributed only to language users. The best reason for adopting the contrary order of explanation, for treating the sort of nonlinguistic belief that is manifested in behavior that can be construed instrumentally as fundamental—seeking thereby to explain the sort of belief that is essentially and not merely accidentally expressible in speech acts—is that it is clear that there were nonlinguistic animals before there were linguistic ones, and the latter did not arise by magic. If the instrumental sense of belief could be made sense of antecedently, and the linguistic sense explained in terms of it, the prospects for explaining how linguistic practice could come into the world in the first place would be bright. This is a laudable aspiration, and it may seem perverse to spurn it. Yet it is a consequence of the account of propositional contents to be offered here that they can be made sense of only in the context of *linguistic* social practices, which have as their core the interpersonal communication of information by assertions. Likewise, rational agency, on which instrumental behavior is modeled, depends essentially on specifically linguistic practices, including asserting. It follows that simple, nonlinguistic, instrumental intentionality can *not* be made fully intelligible apart from consideration of the linguistic practices that make available to the interpreter (but *not* to the interpreted animal) a grasp of the propositional contents attributed in such intentional interpretations.

A theory with such a consequence obviously involves a collateral commitment to show that the conceptual priority of the linguistic sense of belief need not make mysterious the advent of linguistic practices from the capacities of hitherto nonlinguistic creatures. The story to be told here assumes only that suitable social creatures can learn to distinguish in their practice between performances that are treated as correct by their fellows (itself a responsive discrimination) and those that are not. In accord with the pride of place being granted to the linguistic sense of belief, no appeal will be made to instrumental rationality on the part of fledgling linguistic practitioners. The primary explanatory target is what it is to grasp a propositional content and to have and to attribute to others states and performances with such contents—in the sense of explaining what doing the trick consists in, what would count as doing it, rather than how it is done by creatures wired up as we are. Thus no attempt will be made to show how the linguistic enterprise might have gotten off the ground in the first place. But it should be clear at each stage in the account that the abilities attributed to linguistic practition-

ers are not magical, mysterious, or extraordinary. They are compounded out of reliable dispositions to respond differentially to linguistic and nonlinguistic stimuli. Nothing more is required to get into the game of giving and asking for reasons—though to say this is not to say that an interpretation of a community as engaged in such practices can be paraphrased in a vocabulary that is limited to descriptions of such dispositions. Norms are not just regularities, though to be properly understood as subject to them, and even as instituting them by one's conduct (along with that of one's fellows), no more need be required than a capacity to conform to regularities.

If the strategy Field proposes is pursued by looking at the use of the sentences of a public language to perform communicative speech acts, the two subproblems into which he divides the problem of how to understand the attribution of intentionally contentful states correspond to a fundamental pragmatic question and a fundamental semantic question. The first concerns what it is for the utterance of a linguistic expression to have the pragmatic significance of an assertion. This can be rephrased as the question of what it is to use a sentence to make an assertion, provided it is remembered both (1) that it cannot be assumed that *sentences* can be distinguished from other linguistic expressions in advance of saying what it is to use an expression to make an assertion and (2) that *use* should not be assumed to involve a deliberate instrumental exercise of rational agency. Chapter 1 recommended a broadly *normative* approach to the pragmatic question; Chapter 2 recommended a broadly *inferential* approach to the semantic question. It is the task of the rest of this chapter to weave these approaches together into an account of discursive practice—comprising implicitly normative, inferentially articulated statuses, attitudes, and performances. It is the role they play in discursive practice that confers on them objectively representational content, in the most basic case objectively representational *propositional* content. The capacity of practice to confer such contents depends essentially, it will be argued, on its being not only *social* practice but *linguistic* social practice, in that at its core is *communication*, specifically by practitioners' interpretation of each others' *assertions*.

Addressing the pragmatic limb of the two-stage explanatory strategy by appealing to the speech act of assertion yields a further subdivision of issues. As just indicated, it requires an account of what it is for a performance to have the force or pragmatic significance of an assertion, for it to function as an assertion in the game of giving and asking for reasons. Such an account, however, would not by itself yield a suitable reading of Field's *belief\** relation between a sentence and a potential subject of intentional interpretation or attribution of intentional states. For however tightly the two might be linked, there is a substantial difference between believing that *p* and claiming that *p*. (Commitment to a suitable resolution of the semantic subproblem entails a commitment to the eventual appropriateness of this sort of paraphrase in terms of propositional contents rather than sentences.) No sort of intentional

state (or normative status) that might be reconstructed in terms of assertion will provide a suitable analog for belief unless it preserves this contrast by leaving room for the possibility of being in the relevant intentional state without producing the corresponding assertional performance, and for producing the assertional performance without being in the corresponding intentional state. So besides an account of what the assertional significance of a speech act consists in, an account is required also of how attribution of such a speech act is to be understood to be related to attribution of the intentional state it expresses. Such an account might appeal to dispositions—for instance treating being in the state as being *disposed*, under appropriate conditions, to perform the speech act. Or it might appeal to norms—for instance treating the performance of the speech act as involving a *commitment* (which might or might not be fulfilled) involving the state. Or the account might involve both.

## II. DEONTIC STATUS AND DEONTIC ATTITUDES

### 1. *Doxastic Commitments*

The leading idea of the account to be presented here is that belief can be modeled on the kind of inferentially articulated *commitment* that is undertaken or acknowledged by making an assertion. These may be called *doxastic* or *assertional* commitments. This is the basic kind of *discursive* commitment. The strategy is to describe a simplified system of social practices in which something can be taken or treated *as* (having the significance of) an assertion—the acknowledging of commitment to an assertible content. “Assertible content” is what Frege’s “judgeable content”<sup>17</sup> becomes from the point of view of an explanatory commitment to understand judgments in the first instance as what is expressed by assertions. Specifically *propositional* contents (believables) are accordingly to be picked out by the pragmatic property of being assertible. Likewise, what is uttered or inscribed in producing an assertional performance is thereby recognizable as a declarative *sentence*. The role of *propositional* contents marks off *discursive* practice, and the role of *sentential* expressions of such contents is distinctive of *linguistic* social practice. In this way, everything comes down to being able to say what it is for what practitioners are doing to deserve to count as adopting a practical attitude of acknowledging the *assertional* significance of a performance: taking or treating it *as* an assertion. It is in terms of such attitudes that the pragmatic significance of assertional speech acts, the normative status of assertional commitments, and the possession or expression of propositional semantic contents are to be understood.

That the contents conferred by those practices are recognizable as *discursive* or *conceptual* contents (the genus of which propositional contents form the most basic species) depends on their *inferential* articulation and relation

to each other. The practices that institute the significance characteristic of assertional performances and the status characteristic of assertional commitments must be inferential practices. Asserting cannot be understood apart from inferring. So one fundamental question is, What makes something that is done according to a practice—for instance the production of a performance or the acquisition of a status—deserve to count as *inferring*? The answer developed here is that inferring is to be distinguished as a certain kind of move in the game of giving and asking for reasons. To say this is to say that inferring should be understood in an interpersonal context, as an aspect of an essentially *social* practice of communication.

The contentfulness of the states attributed as part of a simple intentional interpretation of an individual consists in a sort of inferential articulation that is not intelligible solely in terms of the role those states play in practical reasoning—if practical reasoning is conceived of as restricted to the sort of intrapersonal instrumental deliberation implicitly imputed by such interpretation. The explanation of behavior according to the model of rational agency depends on treating attributed intentional states as having propositional contents, which involve objective truth conditions. But there is, it will be argued, no pattern of moves a single individual might make that would qualify that individual's states as inferentially articulated in this sense. The inferential practice (including practical reasoning) that confers contents of this kind comprises not only first-person reasoning but also third-person attributions and assessments of it—and both aspects are essential to it. Deliberation is the internalization of the interpersonal, communicative practice of giving reasons to and asking reasons of others, just as judgment is the internalization of a public process of assertion. Inferring cannot be understood apart from asserting. To say this is to say that inferring should be understood as an aspect of an essentially *linguistic* practice. The practice of giving and asking for reasons must be conceived as including assertion because, although there are other kinds of performances besides assertings that can stand in need of reasons—indeed for a performance to be an action just is for it to be something it is appropriate to demand a reason for—to offer a reason is always to make an assertion.

The rest of this chapter is given over to developing a model of assertional and inferential practice. What is described is not our actual practice but an artificial idealization of it. Simplified and schematic though the model may be, it should nonetheless be recognizable as a version of what we do. The model is intended to serve as the core of a layered account of linguistic practice. Where our practice diverges from that specified in the model, those divergences should be explicable as late-coming additions to or modifications of the underlying practice. For instance, the model appeals only to semantic inferences, that is, inferences involving what is *claimed*. Pragmatic inferences such as Gricean implicatures have to do rather with the antecedents and consequents of the performance of *claiming* it. These pragmatic inferen-

tial practices form a shell around the more basic semantic ones, which they presuppose. The critical criterion of adequacy the model answers to is that the core linguistic practices it specifies be *sufficient* to confer propositional and other conceptual contents on the expressions, performances, and deontic statuses that play appropriate roles in those practices. It is also claimed, however, that the fundamental structural features of the model provide *necessary* conditions for the conferral of such contents. So there is a subsidiary commitment to the effect that sophisticated linguistic practices of the sort not addressed by the model are ultimately intelligible only by showing how they could develop out of the sort of practices the model does specify.

## 2. Commitment and Entitlement

At the core of discursive practice is the game of giving and asking for reasons. Chapter 1 sought to motivate the claim that discursive practice is implicitly *normative*; it essentially includes assessments of moves as correct or incorrect, appropriate or inappropriate. The institution of these proprieties by practical assessments on the part of the practitioners is the ultimate source of the meanings of the noises and marks they make, and of the other things they do.<sup>18</sup> As the term is used here, to talk of practices is to talk of *proprieties* of performance, rather than of regularities; it is to prescribe rather than describe. The general notion of proprieties of practice in terms of which the discussion of implicit norms has been conducted up to this point, however, does not cut fine enough to pick out what is distinctive of discursive norms. For that purpose the pragmatics Dummett suggests—which specifies the significance of linguistic expressions (and implicitly of speech acts and alterations of intentional states) in terms of circumstances of *appropriate* application and *appropriate* consequences of such application—must be further refined. Different sorts of propriety must be acknowledged.

The fundamental normative concept required is the notion of *commitment*. Being committed is a normative status—more specifically a *deontic* status. The project of the central sections of this chapter is to introduce a notion of *discursive commitment* as a species of deontic status that can do much of the explanatory theoretical work that is normally assigned to the notion of intentional state. But deontic statuses come in two flavors. Coordinate with the notion of *commitment* is that of *entitlement*. Doing what one is committed to do is appropriate in one sense, while doing what one is entitled to do is appropriate in another. The model of linguistic practice described here elaborates on the Dummettian bipartite pragmatics by distinguishing on the side of consequences, for instance, what a particular speech act commits one to from what it entitles one to. This permits a finer-grained specification of functional roles in linguistic practice than does using a single-sorted notion of propriety of circumstances and consequences of performance.

Commitment and entitlement correspond to the traditional deontic primitives of obligation and permission. Those traditional terms are avoided here because of the stigmata they contain betraying their origin in a picture of norms as resulting exclusively from the commands or edicts of a superior, who lays an obligation on or offers permission to a subordinate. Framed this way, the question of what one is obliged or permitted to do can slip insensibly into the question of who has a right to impose those statuses (as it does explicitly for Pufendorf, for instance). The picture presented here does not depend on a hierarchy of authority. The concepts of obligation and permission, as of duties and rights, can be reconstructed in terms of commitment and entitlement as they will be construed here.

Another way in which the treatment here of the deontic primitives of commitment and entitlement differs from that usually accorded to obligation and permission concerns the relation between them. It has been traditional to acknowledge the relations between these deontic modalities by defining one in terms of another: being permitted to do something is to be rendered as not being obliged not to do it, or being obliged as not being permitted not to. It does make sense to think of being committed to do something as not being entitled not to do it, but within the order of explanation pursued here it would be a fundamental mistake to try to exploit this relation to define one deontic status in terms of the other. Doing so requires taking a formal notion of negation for granted. The strategy employed here is rather to use the relation between commitment and entitlement (which are not defined in terms of this relation) to get a grip on a *material* notion of negation, or better, incompatibility. Two claims are incompatible with each other if commitment to one precludes entitlement to the other. One of the prime advantages normative-functional analyses of the notion of intentional states have over causal-functional analyses is that rendering the phenomenon of <sup>s</sup>inconsistent<sup>s</sup> beliefs<sup>19</sup> as incompatible commitments makes it intelligible in a way not available to causal accounts.

It may also be remarked, in a preliminary fashion, that supposing that sense can be made of the underlying deontic statuses of commitment and entitlement, the notion of material incompatibility of commitments they give rise to leads in a straightforward way to a notion of the *contents* of such commitments. For the content of a commitment can for many purposes be represented by the theorist as the set of commitments that are incompatible with it. For instance, a kind of entailment relation is induced on commitments by inclusion relations among such sets of incompatibles. The commitment *p* incompatibility-entails the commitment *q* just in case everything incompatible with *q* is incompatible with *p*. Thus "Wulf is a dog" incompatibility-entails "Wulf is a mammal," since everything incompatible with his being a mammal is incompatible with his being a dog. The notion of material incompatibility that is made available by *not* defining commitment and entitlement in terms of one another accordingly provides a route from



the pragmatics that deals with these deontic statuses to the semantics that deals with their contents. How this hint might be exploited will become clearer presently.

### 3. *Attributing and Undertaking Commitments*

Deontic statuses of the sort to be considered here are creatures of practical attitudes. There were no commitments before people started treating each other as committed; they are not part of the natural furniture of the world. Rather they are social statuses, instituted by individuals attributing such statuses to each other, recognizing or acknowledging those statuses. Considered purely as a natural occurrence, the signing of a contract is just the motion of a hand and the deposition of ink on paper. It is the undertaking of a commitment only because of the significance that performance is taken to have by those who attribute or acknowledge such a commitment, by those who take or treat that performance as committing the signatories to further performances of various kinds. Similarly for entitlements. A license, such as a ticket, entitles one to do something. Apart from practices of treating people as entitled or not, though, there is just what is actually done. The natural world we consider when bracketing the influence of such social practices contains no distinction between performances one was entitled to and those one was not.

It will be useful to see how this basic vocabulary can be used to discuss the authority and responsibility involved in familiar sorts of deontic statuses. The way in which such statuses can be instituted by practical attitudes can be illustrated by artificially simplified versions of some fundamental normative practices. Authority may be considered first, apart from responsibility. A license, invitation, or entrance ticket entitles or authorizes one to do something one was otherwise not entitled to do. It is always a license in the eyes of someone, for example a ticket-taker or doorman. These "consumers" of licenses (along with the others whose attitudes make the practice into a going concern) constitute them by attributing the authority they thereby come to possess. They do so by treating the authorized one as entitled to a performance. It is not appropriate to enter unless one is authorized by a ticket acquired in the appropriate way. Being given a ticket by the ticket-giver is being authorized or entitled to enter, because and insofar as the ticket-taker will not treat entry as appropriate unless so authorized. In the simplest case, the ticket-taker is the attributor of authority, the one who recognizes or acknowledges it and who by taking the ticket as authorizing, makes it authorizing, so instituting the entitlement.

Practices of this sort can be described in purely responsive terms for prelinguistic communities. The entitlement given and recognized in these practices has a content for an attributor insofar as that attributor practically partitions the space of possible performances into those that have been

authorized and those that have not, by being disposed to respond differently in the two cases. These sanctioning responses (for instance admitting versus ejecting) and the performances they discriminate (enterings of the theater) can be characterized apart from and antecedent to specification of the practice of conferring and recognizing entitlement defined by their means. For this reason the entitling authority will be said to be *externally* defined. The sanctions applied in taking or treating someone as entitled can be specified in nonnormative terms.

The basic structure just considered involves entitlement without commitment, authority without responsibility. A corresponding way into the basic structure of commitments or responsibilities is provided by describing a similarly simplified and artificial version of an actual eighteenth-century British practice. According to this practice, taking "the queen's shilling" from a recruiting officer counts as committing the recipient to military service. A performance of this kind has the same significance that signing a contract would have—in either case one has joined the army and undertaken all the commitments entailed by that change of status. (The official rationale was that some such overt irrevocable nonlinguistic performance was required to do duty for signing a contract, given that those enlisting were largely illiterate. The actual function of the practice was to enable "recruiting" by disguised officers, who frequented taverns and offered what was, unbeknownst to their victims, the queen's shilling, as a gesture of goodwill to those who had drunk up all of their own money. Those who accepted found out the significance of what they had done—the commitment they had undertaken, and so the alteration of their status—only upon awakening from the resulting stupor.) The significance of a commitment is to be understood in terms of the practical attitude of those attributing it, that is in terms of what taking or treating someone *as* committed consists in. In this case, it consists in being liable for punishment by a court-martial if one fails to discharge or fulfill the commitment. The content of the commitment attributed corresponds to the subsequent behavior that would or would not elicit a sanction. Or more precisely, that content can be thought of according to Dummett's model of circumstances and consequences; the particular sanctions (court-martial) are the consequences, and the various failures to perform as a soldier that elicit them are the circumstances.

Two features of this simple commitment structure are worthy of note. First, for someone to *undertake* a commitment, according to this story, is to do something that makes it appropriate to *attribute* the commitment to that individual. That the performance of accepting the coin has the significance of altering the *status* of the one whose performance it is consists in the change it brings about in what *attitudes* are in order. It is by reference to the attitudes of others toward the deontic status (attributing a commitment) that the attitude of the one whose status is in question (acknowledging or undertaking a commitment) is to be understood. So all that is required to make

sense of the normative significance of the performance as an undertaking of commitment is an account of what it is to take or treat someone as committed to do something. The possibility of sanctioning failure to perform appropriately—that is, as one is (thereby) taken to be committed to do—offers a way of construing this fundamental practical deontic attitude.

Second, the basic notion of responsibility or commitment that is introduced by consideration of this simple practice can be understood in terms of the notion of authority or entitlement already discussed. For undertaking a commitment can be understood as authorizing, licensing, or entitling those who attribute that commitment to sanction nonperformance. Such sanctioning would be inappropriate (and so itself subject to sanction) unless it had been authorized by the undertaking of a commitment. One may not court-martial someone who has never joined the service. Thought of this way, the effect of undertaking a commitment is not a matter of in fact eliciting punishment if one does not fulfill the commitment but rather of making such punishment *appropriate*. It is not a matter of the actual conditional dispositions to sanction of those who attribute the commitment but a matter of the conditional normative status of such sanctions. What is being considered is a slightly more sophisticated practice, in which the significance of taking the queen's shilling (the consequences of undertaking a commitment and thereby acquiring a new deontic status) is itself defined in terms of deontic statuses—in particular, of *entitlement* to punish.

The significance of that deontic status (entitlement to punish) might itself be defined in normatively *external* terms; those who attempt to court-martial someone who has not committed himself to the service (and so entitled their superiors to hold such a court) are taken out and shot. Or the difference that entitlement makes might itself be cashed out only in normative terms; those who attempt to court-martial someone who has not committed himself to the service (and so entitled their superiors to hold such a court) thereby *license* or *entitle* their superiors to have them taken out and shot, make it *appropriate* or *correct* for them to be taken out and shot. Normatively *internal* definitions of the significance of changes of deontic status, which specify the consequences of such changes in terms of further changes of deontic status, link various statuses and attitudes into *systems* of interdefined practices.

#### 4. *The Example of Promising*

These points can be illustrated by considering an idealized version of a more familiar sort of practice of undertaking commitments, namely *promising*. Promising is another way of undertaking a task responsibility: committing oneself to perform in a certain way. A special defining feature of promising performances is that they involve specifying what one is committed to do by explicitly *saying* it—*describing* the performances that would

count as fulfilling the commitment. This feature of promising will not be officially intelligible in terms of the theory presented here until the practice of assertion has been introduced. For the promise offers a linguistic characterization of a performance, and that characterization (or a grammatical transform of it) must turn out to be assertible, on pain of the promisor's having failed to fulfill the commitment undertaken. Only one who claims that MacArthur returned will take it that MacArthur's promise to return was fulfilled. Furthermore, promises are typically made *to* someone. The promisees are the ones who are entitled to hold the promiser responsible.

A social-practical description of promising displays how the deontic attitudes of undertaking and attributing the commitment that is a promise are two sides of one coin, aspects of a single social practice. Promising (like taking the queen's shilling) can be understood as consisting in a dimension of responsibility and a coordinate dimension of authority. The authority of the licensing, in the eyes of the attributors of the promise-commitment, is an entitlement on the part of others to rely (as perhaps they were not previously entitled to do) upon the promised performance. The responsibility consists in conditionally entitling others to sanction the committed one. A commitment to perform includes a license to those who attribute the commitment to hold the committed one responsible for nonperformance. To be entitled to hold responsible is to be conditionally entitled to sanction, in case of nonperformance. In the simplest sort of promising practice, both the promised performances and the conditionally authorized sanctions can be specified antecedently (that is, defined externally to the definition of the practice that then ties them together). The promised performance might be washing the dinner dishes, and the conventionally authorized sanction for nonperformance might be being beaten with sticks.

Commitments to perform that are externally defined by practical attitudes manifested in fulfilling performances, and sanctions that are antecedently specifiable in nonnormative terms, share some features with assertional commitments. As already indicated, however, this category must be developed in the direction of a system of practices governing interlocking internally specified significances, defined in terms of fulfilling performances and sanctions. The toy promising practice just mentioned can be developed further by introducing an internally defined sanction for failure to perform as promised. Instead of responding to the failed promiser by beatings or refusals of entry to the theater, attributors of commitments to perform might rather withhold their recognition of that individual's entitlement to undertake such commitments. If the commitment is not fulfilled, the cost is that those who attribute both the commitment and the failure will not (or will become less disposed to) recognize the authorization (making entitled) of reliance on performances, that is, will not recognize the promises of the failed commitment undertaker as counting for anything. Undertaking the commitment is still doing something that conventionally has the significance of entitling others to attribute the commitment, that is both to attribute the commit-

ment and to hold the undertaking performer to it. One who succeeds in making a promise still authorizes others (makes them entitled) to rely on one's future performance, to hold one responsible for a failure to perform according to one's commitment.

These are both authority and responsibility, adding up to commitment. The commitment is undertaken by the one who authorizes or licenses, and undertaking is always licensing the attribution of—that is the holding of the undertaker responsible for—the commitment. Responsibility is of authorizer-undertaker to attributor. The difference is that in the sort of promising being considered, uttering certain words is not always sufficient to undertake the commitment that defines making a promise. For as a sanction consequential on previous failures to keep commitments undertaken by promises, attributors may withhold their recognition of the undertaker's entitlement to issue licenses of that sort, refuse to recognize as successful attempts to undertake such commitments. Unlike the queen's shilling case, no performance compels attribution of authority and responsibility—that is, commitment. Such a sanction is defined internally to the practice being considered, since apart from the practice of promising, one cannot specify what the sanction is.<sup>20</sup>

Part of the definition of what it is to undertake a commitment to perform by promising is the significance of this status for the deontic attitudes of others—the practical interpretation that attributors are entitled to place on the undertaking—namely their right to rely on future performance. The second sort of responsibility undertaken is not a task responsibility (as the commitment to perform could be called). It is a becoming liable to be held responsible (taken to be responsible) for failure to perform as one promised. What the promiser entitles others to do, in this simple practice, is to withdraw their recognition (taking) of one's entitlement to issue a license, of one's authority. When responsibility of this sort is added to authority of the invitation or ticket-giving variety, besides the conditional task responsibility, there must also be the special sort of authority involved in undertaking responsibility. This is authorizing others to hold the undertaker responsible. It is a license to do something, conditional upon the undertaker's failure to fulfill his or her task responsibility once its conditions are satisfied. A distinct and important species of such authority to hold responsible arises when the licensed consequences of failure to perform consist in withholding recognition of entitlement to undertake further responsibilities of the same form. Both promising, as reconstructed here, and asserting, as discussed below, exhibit this special commitment-entitlement structure.

##### 5. *Social Practice: Deontic Statuses and Deontic Attitudes*

The discussion of these simplified examples introduced a number of points. The sort of implicitly normative practice of which language use is a paradigm is to be discussed in terms of two sorts of *deontic status*, namely

commitment and entitlement. The notion of normative status, and of the significance of performances that alter normative status, is in turn to be understood in terms of the practical deontic *attitude* of taking or treating someone as committed or entitled. This is in the first instance *attributing* a commitment or entitlement. Adopting this practical attitude can be explained, to begin with, as consisting in the disposition or willingness to impose *sanctions*. (Later, in more sophisticated practices, *entitlement* to such a response, or its *propriety*, is at issue.) Attributors of these statuses may punish those who act in ways they are not (taken to be) entitled to act, and those who do not act in ways they are (taken to be) committed to act. What counts as punishment may (according to the one who interprets a community as exhibiting practices of this sort) be specifiable in nonnormative terms, such as causing pain or otherwise negatively reinforcing the punished behavior. Or what counts as punishment with respect to a particular practice may be specifiable only in normative terms, by appeal to alterations in deontic status or attitude. A performance expresses the practical attitude or has the significance of an *undertaking* of a commitment in case it *entitles* others to attribute that commitment. So there are two sorts of practical deontic attitudes that can be adopted toward commitments: attributing them (to others) and acknowledging or undertaking them (oneself). Of these, attributing is fundamental.

Here, then, is a way of thinking about implicitly normative social practices. Social practices are games in which each participant exhibits various deontic statuses—that is, commitments and entitlements—and each practically significant performance alters those statuses in some way. The significance of the performance is how it alters the deontic statuses of the practitioners. Looking at the practices a little more closely involves cashing out the talk of deontic *statuses* by translating it into talk of deontic *attitudes*. Practitioners take or treat themselves and others as having various commitments and entitlements. They *keep score* on deontic statuses by attributing those statuses to others and undertaking them themselves. The significance of a performance is the difference it makes in the deontic score—that is, the way in which it changes what commitments and entitlements the practitioners, including the performer, attribute to each other and acquire, acknowledge, or undertake themselves. The significance of taking the queen's shilling lies in its being an undertaking of a commitment on the part of the recipient, altering the attributions of commitment by those who appreciate the significance of the performance. It entitles other authorities—those who according to the antecedent score already had undertaken various commitments or duties and entitlements or sorts of authority, those who therefore play a certain role or hold a certain office in the system of practices in question—to punish the performer in particular ways under particular circumstances. The normative significances of performances and the deontic states of performers are instituted by the practice that consists in keeping score by adopting attitudes of attributing and acknowledging them.<sup>21</sup>

### III. ASSERTING AND INFERRING

#### 1. *Linguistic Practice: Assertion and Inference*

The discussion of the significance of performances as altering the deontic attitudes that keep track of normative statuses has so far addressed implicitly normative social practices—whether or not they are specifically *linguistic* (and so more generally *discursive*) practices. What distinguishes the latter sort is the *inferential* articulation of the normative significances they involve—and so their conferral of specifically *conceptual* content on the states, attitudes, performances, and expressions they govern. The challenge is to show how these two approaches (normative pragmatics modeled on deontic scorekeeping, and inferential semantics) can be combined into a single story about social practices of treating speech acts as having the significance of assertions.

Describing practices sufficient to institute such a significance is the way to fill in the notion of assertional commitment. Such an account provides an answer to the question, What is it that we are *doing* when we assert, claim, or declare something? The general answer is that we are undertaking a certain kind of *commitment*. Saying specifically *what* kind is explaining what structure must be exhibited by the practices a community is interpreted as engaging in for that interpretation to be recognizable as taking the practitioners to be keeping score for themselves and each other in virtue of the alterations of their practical deontic attitudes of attributing and undertaking assertional commitments and their corresponding entitlements.

The key to seeing how the scorekeeping model of deontic social practices can be used to make sense of asserting is Sellars's notion of a "game of giving and asking for reasons." The idea is that assertings (performances that are overt undertakings of assertional commitments) are in the fundamental case what reasons are asked for, and what giving a reason always consists in. The kind of commitment that a claim of the assertional sort is an expression of is something that can stand in need of (and so be liable to the demand for) a reason; and it is something that can be offered as a reason. This is the principle motivating the present strategy for discriminating assertional commitments from other species of commitment. Other things besides assertional commitments involve liability to demands for justification or other demonstration of entitlement—for instance, the practical commitments involved in *actions*. Other things besides assertional commitments can entitle interlocutors to assertional commitments—for instance *reliability* in the responsive acquisition of assertional commitments of a certain kind. For being a reliable reporter of currently visible red things who responsively acquires a disposition to claim that there is something red in the vicinity may entitle someone to that commitment.<sup>22</sup> But only assertional commitments stand in *both* these relations.

That claims play the dual role of justifier and subject of demand for justification is a necessary condition of their kind properly being called

*assertional* commitments. It is here employed as well as part of a sufficient condition, in an idealized artificial practice constructed to model this central aspect of the use of natural language. Specifically *linguistic* practices are distinguished as just the social practices according to which some performances have the significance of undertakings of assertional commitment (in virtue of their role in giving and asking for reasons); *declarative sentences* are picked out as the expressions uttered or inscribed in such assertional performances. What is expressed by such performances and determine the particular features of their significance within the assertional genus count thereby as *propositional* contents.

The idea exploited here, then, is that *assertions* are fundamentally fodder for *inferences*. Uttering a sentence with assertional force or significance is putting it forward *as* a potential reason. Asserting is giving reasons—not necessarily reasons addressed to some particular question or issue, or to a particular individual, but making claims whose availability as reasons for others is essential to their assertional force. Assertions are essentially *fit* to be reasons. The function of assertion is making sentences available for use as premises in inferences. For performances to play this role or have this significance requires that assertional endorsement of or commitment to something entitles or obliges one to other endorsements. The pragmatic significance of assertional commitments and entitlements to such commitments consists in the ways in which they are *heritable*; their heritability is the form taken by the *inferential* articulation in virtue of which they count as semantically contentful.

## 2. Three Dimensions of Inferential Articulation

The basic model of the inferential practices that institute assertional significance—and thereby confer propositional contents on states, attitudes, and performances playing suitable roles in those practices—is defined by a structure that must be understood in terms of the interaction of three different dimensions. First, there are two different sorts of deontic status involved: *commitments*, and *entitlements* to commitments. Inheritance of commitment (being committed to one claim as a consequence of commitment to another) is what will be called a *committive*, or commitment-preserving, inferential relation. Deductive, logically good inferences exploit relations of this genus. But so do materially good inferences, such as inferences of the form: *A* is to the West of *B*, so *B* is to the East of *A*; This monochromatic patch is green, so it is not red; Thunder now, so lightning earlier. Anyone committed to the premises of such inferences is committed thereby to the conclusions.

Inheritance of entitlement (being entitled to one claim as a consequence of entitlement to another) is what will be called a *permissive*, or entitlement-preserving inferential relation. Inductive empirical inferences exploit rela-



tions of this genus. The premises of these inferences entitle one to commitment to their conclusions (in the absence of countervailing evidence) but do not compel such commitment. For the possibility of entitlement to commitments incompatible with the conclusion is left open. In this way the claim that this is a dry, well-made match can serve as a justification entitling someone to the claim that it will light if struck. But the premise does not commit one to the conclusion, for it is compatible with that premise that the match is at such a low temperature that friction will not succeed in igniting it. The interplay between the two sorts of deontic status is at the center of the model of assertional and inferential practices presented here.

The broadly inferential roles that are identified with propositional contents involve not only commitment- and entitlement-preserving inferential connections among such contents but also relations of *incompatibility*. To say that two claims have (materially) incompatible contents is to say that commitment to one precludes entitlement to the other.<sup>23</sup> No candidate notion can count as a construal of the sort of propositional content we take assertions, judgments, and beliefs to have unless it underwrites incompatibility relations among them. (That possible worlds, for instance, must be understood as corresponding to maximal sets of *compatible* propositions is acknowledged both by those who want to exploit that principle to define propositions and their material compatibility in terms of possible worlds and by those who would reverse that order of explanation.) The explanatory strategy adopted here is to begin with practices that institute deontic statuses of commitment and entitlement and then to show how those practices thereby confer specifically propositional conceptual contents on what is assertible—contents recognizable as such in virtue of the deontic inferential and incompatibility relations they stand in.

The second dimension of broadly inferential articulation that is crucial to understanding assertional practice turns on the distinction between the *concomitant* and the *communicative* inheritance of deontic statuses. This is the *social* difference between *intrapersonal* and *interpersonal* uses of a claim as a premise. Undertaking a commitment or acquiring an entitlement has consequences for the one whose statuses those are. One commitment carries with it other concomitant commitments as consequences. Its consequences are those that it entails according to the commitment-preserving inferential relations that its content stands in to other contents of possible commitments. Similarly, being entitled to a commitment can entitle one to others, which stand to it in suitable permissive or entitlement-preserving inferential relations. Again, the definition of incompatibility of contents in terms of commitment and entitlement means that acquiring a commitment may have as a consequence the loss of entitlement to concomitant commitments one was heretofore entitled to.

But these intrapersonal inferential consequences of changes in deontic status do not exhaust the significance of assertional performances. Such

performances also have a significance for interpersonal *communication*. Putting a sentence forward in the public arena *as true* is something *one* interlocutor can do to make that sentence available for *others* to use in making further assertions. Acknowledging the undertaking of an assertional commitment has the *social* consequence of licensing or entitling others to *attribute* that commitment. The adoption of that deontic attitude on the part of the audience in turn has consequences for the commitments the audience is entitled to undertake. Putting a claim forward *as true* is putting it forward as one that it is appropriate for others to *take true*, that is to endorse themselves. Assertion that is communicatively successful in the sense that what is put forward as true by a speaker is taken as true by the audience consists in the interpersonal inheritance of commitment.

The third dimension of broadly inferential articulation that is crucial to understanding assertional practice is that in which discursive *authority* is linked to and dependent upon a corresponding *responsibility*.<sup>24</sup> In uttering a sentence assertively, the claim one makes involves an *endorsement*. One aspect of this sort of endorsement was indicated above in a preliminary fashion in terms of the function of an asserting as licensing or *authorizing* further assertions (and eventually, actions—but consideration of practical rather than doxastic commitments is postponed until the next section). But unless some independent grasp is offered of the status or significance that must be bestowed on a performance for it to count as an asserting, invoking the inferential warranting of further assertions merely traces out a rather small circle. It is the second aspect of endorsement, of the sort of *responsibility* involved in assertional commitment, that permits a larger horizon. Understanding that aspect requires putting together the distinction between the deontic statuses of commitment and entitlement, on the one hand, and of intra- and interpersonal inference, on the other. Putting forward a sentence as true or as information—that is, asserting it—has been glossed as putting it forward as fit to be a reason for other assertions, making it available as a premise from which others can be inferred. This means that others can inherit *entitlement* to an assertional commitment from the one who makes an assertion and thereby licenses or warrants its reassertion and the assertion of what follows from it. To understand this warranting function, the heritability of entitlement, requires understanding the social significance of the distinction between warranted and unwarranted assertional commitments.

Ordinarily the relation of an authorizing event to the performances it licenses requires at least that in the context of that event, performances become socially appropriate that otherwise would not be. For example purchasing a ticket entitles one to take a seat in the theater, which it would be inappropriate to do without the ticket. This observation presents a dilemma. If asserting a sentence is not a performance requiring prior authorization, then it seems one cannot understand the function of assertion as inferentially licensing further assertions. If, however, asserting is a performance requiring

authorization, how does one become entitled to the original licensing assertion? Talk of *inheritance* of entitlement makes sense only in an explanatory context that includes a story about the significance of *possession* of entitlement. It is this question that is addressed by an account of the dimension of *responsibility* characteristic of asserting. In asserting a claim, one not only authorizes further assertions (for oneself and for others) but undertakes a responsibility, for one commits oneself to being able to vindicate the original claim by showing that one is entitled to make it. Others cannot inherit an entitlement that the asserter does not possess. Overtly acknowledging or undertaking a doxastic commitment by issuing an assertional performance can warrant further commitments, whether by the asserter or by the audience, only if that warranting commitment itself is one the asserter is entitled to. Only assertions one is entitled to make can serve to entitle anyone to their inferential consequences.

The function of asserting as the giving of reasons is intelligible only as part of a practice in which reasons can be asked for or required. That some performances admit or stand in need of reasons is presupposed by the practice of offering them. Many kinds of performances are subject to demands for or explanation according to reasons. The two fundamental sorts discussed here are intentional actions that are not speech acts and assertions themselves. Both actions and assertions—overt performances corresponding to practical and to doxastic commitments respectively—are essential and ineliminable aspects of discursive practice as here construed. Nonetheless, the significance of assertional performances can be filled in to a considerable extent before it is necessary to look at the role of assertions as reasons for anything other than more assertions.

The converse is not the case. Actions just are performances for which it is appropriate to offer reasons, and offering a reason is making an assertion. So actions are not intelligible as such except in a context that includes assertional giving of reasons. Where intentional explanations are offered of the behavior of nonlinguistic creatures (those that are not understood as interpreters of others), the reasons are offered, the assertions are made, by the interpreter of the simple intentional system, who seeks to make its behavior intelligible by treating it *as if* it could act according to reasons it offers itself. That is why what is attributed in such interpretations is *derivative* intentionality. Assertions play both roles; reasons can be offered for them, and they can be offered as reasons. Actions play only the first role in the game of giving and asking for reasons. Thus a description of that game that has a certain autonomy can be offered if, to begin with, attention is restricted to assertions alone, and as a result asserting has a certain explanatory priority over action. Although for this reason the discussion of actions and the practical commitments they express is postponed until the next chapter, it remains that asserting is a *doing*, and the responsibility it involves should also be understood as a responsibility to *do* something.

Besides its specifically linguistic use, 'assert' has a broader normative use according to which one can assert one's authority or one's rights. An important component of this sense of asserting is defending, championing, or justifying. Milton uses the word this way in the famous expression of his intent in *Paradise Lost*: "That to the highth of this great Argument I may assert Eternal Providence, and Justifie the wayes of God to men." This use points to the fact that the sort of commitment involved in linguistic asserting involves the undertaking of a specifically *justificatory* responsibility for what is claimed. In asserting a sentence, one not only licenses further assertions (for others and for oneself) but commits oneself to *justifying* the original claim. The responsibility in question is of the sort Baier calls "task-responsibility"; it requires the performance of a task of some kind for its fulfillment.<sup>25</sup> Specifically, in making a claim, one undertakes the conditional task responsibility to demonstrate one's entitlement to the claim, if that entitlement is brought into question. Justifying the claim when it is queried, giving reasons for it when reasons are asked for, is one way to discharge this obligation. If the commitment can be defended, entitlement to it demonstrated by justifying the claim, then endorsement of it can have genuine authority, an entitlement that can be inherited.

### 3. Assertion as a Doing: Authority and Responsibility

The position maintained here is that *discursive* (in the Kantian sense of concept-mongering) practice can only be *linguistic* practice, and that what distinguishes a practice as specifically linguistic is that within it some performances are accorded the significance of *assertions*. It is only because some performances function as assertions that others deserve to be distinguished as *speech* acts. The class of questions, for instance, is recognizable in virtue of its relation to possible answers, and offering an answer is making an assertion—not in every individual case, but the exceptions (for example, questions answered by orders or by other questions) are themselves intelligible only in terms of assertions. Orders or commands are not just performances that alter the boundaries of what is permissible or obligatory. They are performances that do so specifically by *saying* or *describing* what is and is not appropriate, and this sort of making explicit is parasitic on claiming. Saying "Shut the door!" counts as an order only in the context of a practice that includes judgments, and therefore assertions, that the door is shut or that it is not shut. (The "slab" *Sprachspiel* Wittgenstein describes in the opening paragraphs of the *Investigations* is not in this sense a *language* game—it is a set of practices that include only vocal, but not yet verbal, signals.) In the same way, promises are not just undertakings of responsibility to perform in a certain way. They are performances that undertake such responsibility by *saying* or *describing* explicitly what one undertakes to do. One promises in effect to make a proposition true, and the propositional

contents appealed to can be understood only in connection with practices of saying or describing, of taking-true—in short, of asserting what are, in virtue of the role they play in such assertions, declarative sentences. As it is with these examples, so it is with other speech acts. Asserting is the fundamental speech act, defining the specific difference between linguistic practice and social practices more generally.

A crucial measure according to which a theory of speech acts ought to be assessed, then, is its treatment of what one is *doing* in producing an assertion. This challenge is not always accepted. One prominent theorist defines the assertion of the declarative sentence *p* as “an undertaking to the effect that *p*.”<sup>26</sup> One does not have to subscribe to the pragmatist project of explaining the propositional contents that are asserted in terms of the practices of asserting them in order to find this disappointing. What sort of an undertaking is this? What, exactly, is the effect? The theory being presented here aims to answer just these questions.

In producing assertions, performers are doing two sorts of things. They are first *authorizing* further assertions (and the commitments they express), both concomitant commitments on their part (inferential consequences) and claims on the part of their audience (communicational consequences). In doing so, they become responsible in the sense of answerable for their claims. That is, they are also *undertaking* a specific task *responsibility*, namely the responsibility to show that they are *entitled* to the commitment expressed by their assertions, should that entitlement be brought into question. This is the responsibility to *do* something, and it may be fulfilled for instance by issuing other assertions that *justify* the original claim. The semantic content of the commitment expressed by the performance—that the authority it claims and the justificatory responsibility it undertakes are specifically “to the effect that *p*” (rather than some other *q*)—consists in its specific inferential articulation: what else it commits the asserter to, what commitments it entitles its audience to, what would count as a justification of it, and so on. On this account, then, the pragmatic force or significance characteristic of asserting (and therefore ultimately also the concepts of declarative *sentence*, *propositional* content, and specifically *linguistic* practice) is to be understood in terms of performances with the dual function in the game of giving and asking for reasons of being givings of reasons, and themselves also performances for which reasons can be asked. The conceptual contents expressed by assertional performances are to be explained by appeal to the inferential roles they play in that reason-mongering practice. What is *done* in asserting—the pragmatic *significance* or *effect* of producing an assertional performance—consists in the way in which, by authorizing particular further inferentially related performances and undertaking responsibility to produce yet other inferentially related performances, asserters alter the score interlocutors keep of the deontic statuses (commitments and entitlements) of their fellow practitioners.

The analysis being suggested divides this significance into a component having to do with *authority* and one having to do with *responsibility*. The particular way these components are intertwined defines the sort of pragmatic significance that is being identified with assertional force. The constellation of authority and responsibility characteristic of the assertional significance of speech acts is *socially* articulated. In producing an assertion, one undertakes a responsibility oneself. The authority of that performance (which is conditioned on the responsibility) in turn consists in opening up a new avenue along which those in the audience can fulfill the responsibilities associated with *their* assertions. At the core of assertional practice lie three fundamental ways in which one can demonstrate one's entitlement to a claim and thereby fulfill the responsibility associated with making that claim. Two of these—justifying the content of an assertion and deferring to the authority of an asserter—can be discussed here. The third—invoking one's own authority as a reliable noninferential reporter—is best left for later.

First, as already mentioned, one can demonstrate one's entitlement to a claim by *justifying* it, that is, by giving reasons for it. Giving reasons for a claim always consists in making more claims: asserting premises from which the original claim follows as a conclusion.<sup>27</sup> Interlocutors who accept such a vindication of the commitment—who accept the reasons offered as a justification demonstrating entitlement to the conclusion—thereby implicitly endorse a certain inference. These practical attitudes of taking or treating inferences as correct (distinguishing them from the incorrect ones by responses to attempted justifications) institute inferential proprieties relating the performances of asserters (and the commitments adopted thereby) and so confer contents on them. For it is the practical inferential proprieties acknowledged by such attitudes that make noises and marks mean what they mean. Assertions play a dual role in justification: as justifiers and as justifieds, premises and conclusions. That it plays this dual role, that it is caught up in justificatory inferences both as premise and as conclusion, is what makes it a specifically *propositional* (= assertible, therefore believable) content at all. That it exhibits the particular inferential grounds and consequences that it does is what makes it the particular determinate content that it is—settling, for instance, what information it conveys, the significance that undertaking a commitment with that content would have for what else one is committed and entitled to. Thus the inferential articulation of speech acts depends on this practice of demonstrating entitlement to the commitment acknowledged by the performance of a speech act.

The second way of vindicating a commitment by demonstrating entitlement to it is to appeal to the authority of another asserter. The *communicational* function of assertions is to license others who hear the claim to reassert it. The significance of this license is that it makes available to those who rely on it and reassert the original claim a special way of discharging their responsibility to demonstrate their entitlement to it. They can invoke

the license or authority of the asserter, thus deferring to the interlocutor who communicated the claim and passing along to that other individual any demands for demonstration of entitlement. The authority of an assertion includes an offer to pick up the justificatory check for the reassertions of others. That *A*'s assertion has the social significance of authorizing *B*'s reassertion consists in the appropriateness of *B*'s deferring to *A* the responsibility to demonstrate entitlement to the claim. *B*'s responsibility can be discharged by the invocation of *A*'s authority, upon which *B* exercised the right to rely. The buck is passed to *A*.

So communication does not involve only the sharing of *commitments*—their spreading from one individual to another as the speaker who produces an assertion communicates to, and possibly infects, an audience. It involves also the way that *entitlements* to claims can be inherited by the consumers of an assertion from its producer. In this way the *authority* of an assertional performance consists in part in making available a new way in which those to whom it is communicated can discharge their *responsibility* for demonstrating entitlement to commitments they undertake. Furthermore, assertions can be seen to play a dual role on the side of communication, just as they do on the side of justification. For assertions are on the one hand what is communicated (made available to others), and on the other hand they are what communication is *for*: one interlocutor's claim is fodder for inferences by others to further claims. The audience not only attributes to the one producing an assertional performance commitment to claims entailed (according to commitment-preserving inferences) by the assertion, but it also may undertake commitments and acquire entitlements that are its consequences.

In such inheritance of entitlement by communication, the content of the commitment is preserved intact and merely transferred from one scorekeeper to another. The communicational mechanism for fulfilling the responsibility to demonstrate entitlement appeals to *interpersonal*, *intracontent* inheritance of entitlement to a propositional commitment. By contrast, the justificatory mechanism appeals to *intrapersonal*, *intercontent*<sup>28</sup> inheritance of entitlement to a propositional commitment—since the contents of premises and conclusions will differ in any inference that is nontrivial in the sense of being available to do justificatory work. This combination of the *person-based* authority (invoked by *deferring* to the claim of another) and *content-based* authority (invoked by *justifying* the claim through assertion of other sentences from which the claim to be vindicated can appropriately be inferred) is characteristic of asserting as a doing. This constellation—of commitment and entitlement, of authority and responsibility, and of an inheritance of entitlement to assertional commitments that is *interpersonal* and *intracontent* as well as *intrapersonal* and *intercontent*—constitutes a fundamental substructure of the model of assertional practices presented here.

#### 4. *The Default and Challenge Structure of Entitlement*

More clearly needs to be said about the practices that govern the attribution of entitlement to assertional commitments. The two mechanisms considered so far for demonstrating such entitlement are ways in which entitlement to commitments can be *inherited*. Entitlement to commitment to one claim can be extended to entitlement to another either according to the *inferential* pattern appealed to by justification, in which case that entitlement is inherited by another commitment (to a different claim) undertaken by the same interlocutor, or according to the *communicational* pattern appealed to by deferral, in which case that entitlement is inherited by another commitment (to the same claim) by a different interlocutor. Tracing back an entitlement secured by either of these sorts of inheritance potentially sets off a regress.

The justificatory style of vindication, in which one interlocutor offers premises with different contents as reasons for a claim, threatens a regress on claim contents. At each stage vindication of one commitment may involve appeal to commitments that have not previously been invoked, for which the issue of demonstrating entitlement can arise anew, so the issue is merely put off. Or at some point a circle is closed by appeal to a set of premises whose entitlement has already been brought into question (and put off). Then the argument offered for a claim amounts to something that could be made explicit (eliminating the intervening steps) by a stuttering inference of the form "*p*, therefore *p*," which cannot create entitlement.

The communicational style of vindication, in which one interlocutor appeals to another interlocutor's assertional avowal of a commitment with the same content, threatens a corresponding regress on interlocutors. At each stage vindication of one interlocutor's commitment may involve appeal to the commitment of some interlocutor who has not yet been appealed to, for whom the issue of demonstrating entitlement can arise anew—so the issue is merely put off. Or at some point a circle is closed by appeal to the assertion of some interlocutor whose entitlement has already been brought into question (and put off). Then the deferral that seeks to vindicate the claim amounts to something that could be made explicit (eliminating the intervening steps) by a self-citation of the form "I am relying for my entitlement to *p* on the authority of my own claims that *p*," which cannot create entitlement.

The situation is not fundamentally altered by the fact that tracing back a single entitlement might involve both inferential and communicational appeals—that the chain of inheritance might comprise both justifications and deferrals. These are mechanisms for spreading entitlements, not for originating them; combining the two merely results in more complicated regresses and circles. What gets the process off the ground? What gives these multiplicative mechanisms something to work with in the first place, so that chains of vindication can come to an end? This question arose above in connection



with the authorizing function of assertions and was pursued through the notion of fulfilling the responsibility to vindicate the authorizing commitment, by demonstrating one's entitlement. But looking at those mechanisms raises the same issue all over again.

The worry about a regress of entitlements is recognizably foundationalist. It can be responded to by appealing to the fundamental pragmatic commitment to seeing normative statuses (in this case entitlement) as implicit in the social practices that govern the giving and asking for reasons. Those practices need not be—and the ones that actually confer content on our utterances are not—such that the default entitlement status of a claim or assertional commitment is to be guilty until proven innocent. Even if all of the methods of demonstrating entitlement to a commitment are regressive (that is, depend on the inheritance of entitlement), a grounding problem arises in general only if entitlement is never attributed until and unless it has been demonstrated. If many claims are treated as innocent until proven guilty—taken to be entitled commitments until and unless someone is in a position to raise a legitimate question about them—the global threat of regress dissolves.

One of the lessons we have learned from thinking about hyperbolic Cartesian doubt is that doubts too sometimes need to be justified in order to have the standing to impugn entitlement to doxastic commitments. Which commitments stand in need of vindication (count as defective in the absence of a demonstration of entitlement to them) is itself a matter of social practice—a matter of the practical attitudes adopted toward them by the practitioners. The different circumstances under which various claims are taken or treated as requiring justification (or vindication by deferral) is part of what confers on the sentences that express them the different meanings that they have. It is part of the inferential role they play, in the broad practical sense of that expression, that includes the conditions under which inferential performances of various sorts are appropriate or obligatory. Claims such as “There have been black dogs” and “I have ten fingers” are ones to which interlocutors are treated as *prima facie* entitled. They are not immune to doubt in the form of questions about entitlement, but such questions themselves stand in need of some sort of warrant or justification. Entitlement is, to begin with, a social status that a performance or commitment has within a community.<sup>29</sup> Practices in which that status is attributed only upon actual vindication by appeal to inheritance from other commitments are simply unworkable; nothing recognizable as a game of giving and asking for reasons results if justifications are not permitted to come to an end.

The model presented here has what might be called a *default and challenge structure* of entitlement. Often when a commitment is attributed to an interlocutor, entitlement to it is attributed as well, by default. The *prima facie* status of the commitment as one the interlocutor is entitled to is not permanent or unshakeable; entitlement to an assertional commitment can

be challenged. When it is *appropriately* challenged (when the challenger is *entitled* to the challenge), the effect is to void the inferential and communicative authority of the corresponding assertions (their capacity to transmit entitlement) unless the asserter can vindicate the commitment by demonstrating entitlement to it.

This is what was meant by saying that the broadly justificatory responsibility to vindicate an assertional commitment by demonstrating entitlement to it is a *conditional* task-responsibility. It is conditional on the commitment's being subject to a challenge that itself has, either by default or by demonstration, the status of an entitled performance. Indeed, the simplest way to implement such a feature of the model of asserting is to require that the performances that have the significance of challenging entitlements to assertional commitments themselves be assertions. One then can challenge an assertion only by making an assertion incompatible with it. (Recall that two claims are incompatible just in case commitment to one precludes entitlement to the other.) Then challenges have no privileged status: their entitlement is on the table along with that of what they challenge. Tracing the provenance of the entitlement of a claim through chains of justification and communication is appropriate only where an actual conflict has arisen, where two *prima facie* entitlements conflict. There is no point fixed in advance where demands for justification or demonstration of entitlement come to an end, but there are enough places where such demands *can* end that there need be no *global* threat of debilitating regress.

This is the sort of picture of the practices of giving and asking for reasons that Wittgenstein suggests, but it is recognizable already in Socratic elenchus. In the present context, the proper question is not whether practices that incorporate such a default-and-challenge structure of entitlements are somehow in principle defective in view of some a priori rationalistic criterion of what it is to be *really* entitled to a claim. The proper question is rather, What sort of propositional contents can reason-constituting practices of this sort confer on the scorekeeping attitudes they govern, the deontic statuses they institute, and the performances they acknowledge as having the significance of assertions? The claim eventually to be made is that such practices suffice to confer objectively representational propositional contents on claims, objective truth conditions according to which the correctness of an assertion can depend on how things are with the objects represented by it, to the extent that the entire linguistic community could be wrong in its assessment regarding it.

##### 5. *Internal Sanctions: Doxastic Commitments without Entitlements Lack Authority*

The picture, then, is one in which giving reasons is obligatory only if they have been appropriately asked for. What has the significance of a

challenge (a demand for reasons) is just more assertions, whose entitlements are subject to the same sort of assessments as any others. All are weighed in the same balance. The fundamental concept in terms of which the default-and-challenge structure is adumbrated is the deontic attitude of *attributing* entitlement, of one interlocutor's taking or treating another *as* entitled to a commitment or performance. Now that the background presupposed by an interlocutor's conditional task-responsibility to demonstrate entitlement to the commitment undertaken by an assertional performance has been filled in, it is possible to be a bit clearer about this deontic attitude. The question is, What practical difference does it make whether the asserter is entitled to an assertional commitment? That is, What is the pragmatic significance of the distinction between warranted and unwarranted assertional commitments? What is it about a scorekeeper's treatment of an attributed commitment that makes it appropriate to describe that practical attitude as one of taking the commitment to be one the asserter is not *entitled* to?

Answering this question requires considering all three dimensions of assertion mentioned above: relations between commitments and entitlements, relations between intra- and interpersonal inferential significance, and relations between authority and responsibility. Since the *authority* of an assertion consists in its inferentially licensing or warranting further commitments, and this is a matter of *inheritance* of entitlement, an assertion expressing a commitment that is not taken to be one the performer is entitled to will not be taken to have inferential authority. Although making a claim by asserting a sentence is putting it forward *as* a fit premise for oneself and others to draw conclusions from, it will be accorded that status (its authority recognized) only by those scorekeepers who attribute not only the *commitment* the performance expresses but also an *entitlement* to that commitment. Absent such entitlement, assertion is an attempt to lend what one does not possess. Failure to shoulder the justificatory *responsibility* associated with entitlement to an assertional commitment (supposing it to have been appropriately challenged) renders void its *authority* as an inferential warrant for further commitments. Inferential authority and justificatory responsibility are coordinate and commensurate.

In the ideal Sprachspiel being described, making a claim one is not entitled to (even as a challenge) is a kind of impropriety, the violation of a norm. For a performance to have this sort of status or significance within or according to a set of practices—for this sort of norm to be implicit in or instituted by those practices—requires that the practices include attitudes of taking, treating, or acknowledging performances *as* incorrect in that particular way. Some sort of sanction must be available, with respect to which it can be specified what a practitioner *does* in adopting those practical attitudes. The practical significance of lack of entitlement consists in liability to *punishment* of some kind.

As has already been pointed out, however, such punishment need not

consist in *external* sanctions—responses such as beating with sticks, which are interpretable as punishments (perhaps in virtue of functioning as negative reinforcement in a behavioral-statistical sense) apart from the normative significance they in turn have within the practices in question. One can coherently interpret a community as engaging in practices in which performances are treated as having the significance of promises (or of the more primitive sort of nonlinguistic undertaking of task-responsibility, of which taking the queen's shilling is an example) even if the *only* sanction for failing to perform as one has committed oneself to do is to disqualify oneself from counting in the future as eligible to undertake such commitments. Something like this is what happened to the boy who cried "Wolf." Having several times committed himself to the claim that a wolf was present (thereby licensing and indeed obliging others to draw various conclusions, both practical and theoretical) under circumstances in which he was not entitled by the evident presence of a wolf to undertake such a commitment and to exercise such authority, the boy was punished—his conduct practically acknowledged as inappropriate—by withdrawal of his franchise to have his performances treated as normatively significant.

Unlike the case of the liar who ceases to be believed or of the irresponsible promiser who ceases to be relied upon, however, the internal sanctions constituting the practical recognition of an assertional performance as one the performer is not entitled to do not, on the present model, deal with the significance accorded to *other* performances of the same sort by that individual. Those sanctions have rather to do with the significance assigned to that very performance. Treating the commitment expressed by an assertional performance as one the asserter is not entitled to is treating it as not entitling that interlocutor and the audience of the performance to commitments whose contents follow inferentially from the asserted content. The practical sanction constitutive of the implicit norm governing entitlement to assertional commitments is internal to the system of scorekeeping attitudes the practice comprises. Taking someone to be (or not to be) entitled to a claim has consequences for what deferring performances one acknowledges as in order—and this in turn affects what deferrings one is oneself disposed to produce under various circumstances. But there need be no social pattern of performances and dispositions describable in nonnormative terms that is either necessary or sufficient for the constitution of such deontic attitudes.

#### IV. SCOREKEEPING: PRAGMATIC SIGNIFICANCE AND SEMANTIC CONTENT

##### 1. *Lewis's Version of Scorekeeping in Language Games*

The particular way in which the pragmatic significance of speech acts and deontic statuses is related to their semantic contents can be clarified

by looking more closely at the metaphor of *scorekeeping* by linguistic practitioners. The use made of the notion here is an adaptation of an idea introduced by David Lewis, in his paper "Scorekeeping in a Language Game." He suggests thinking about the rule-governedness of conversation by using some of the concepts appropriate to games that evolve according to a score function. This notion is explained in terms of baseball:

At any stage in a well-run baseball game there is a septuple of numbers  $\langle r_v, r_h, h, i, s, b, o \rangle$  which I shall call the *score* of the game at that stage. We can recite the score as follows: the visiting team has  $r_v$  runs, the home team has  $r_h$  runs, it is the  $h^{\text{th}}$  half ( $h$  being 1 or 2) of the  $i^{\text{th}}$  inning; there are  $s$  strikes,  $b$  balls, and  $o$  outs.<sup>30</sup>

The *constitutive rules* of the game are then of two sorts:

*Specifications of the kinematics of score.* Initially the score is  $\langle 0, 0, 1, 1, 0, 0, 0 \rangle$ . Thereafter, if at time  $t$  the score is  $s$ , and if between time  $t$  and  $t'$  the players behave in manner  $m$ , then at time  $t'$  the score is  $s'$ , where  $s'$  is determined in a certain way by  $s$  and  $m$ .

*Specifications of correct play.* If at time  $t$  the score is  $s$ , and between time  $t$  and time  $t'$  the players behave in manner  $m$ , then the players have behaved incorrectly. (Correctness depends on score: what is correct play after two strikes differs from what is correct play after three.) What is not incorrect play according to these rules is correct.<sup>31</sup>

He then points out that it is possible to use specifications of these sorts to define 'score' and 'correct play', by using the notion of *score function*, which is

the function from game-stages to septuples of numbers that gives the score at every stage. The specifications of the kinematics of score, taken together, tell us that the score function evolves in such-and-such way. We may then simply define the score function as the function which evolves in such-and-such way . . . . Once we have defined the score function, we have thereby defined the score and all its components at any stage. There are two outs at a certain stage of a game, for instance, if and only if the score function assigns to that game-stage a septuple whose seventh component is the number 2.<sup>32</sup>

Correct play is specified in terms of current score and current behavior. Since the required relation between these is codified in the score function, it also defines correct play.

The idea is, then, that the evolution of a linguistic interchange or conversation can be thought of as governed by implicit norms that can be made explicit (by the theorist) in the form of a score function. Here are some of the analogies Lewis points to:

—Like the components of a baseball score, the components of a conversational score at any given stage are abstract entities. They may not be numbers, but they are other set-theoretic constructs . . .

—What play is correct depends on the score . . .

—Score evolves in a more-or-less rule-governed way. There are rules that specify the kinematics of score:

If at time  $t$  the conversational score is  $s$ , and if between time  $t$  and time  $t'$  the course of conversation is  $c$ , then at time  $t'$  the score is  $s'$ , where  $s'$  is determined in a certain way by  $s$  and  $c$ .

Or at least:

. . . then at time  $t'$  the score is some member of the class  $S$  of possible scores, where  $S$  is determined in a certain way by  $s$  and  $c$  . . .

—To the extent that conversational score is determined, given the history of the conversation and the rules that specify its kinematics, these rules can be regarded as constitutive rules akin to definitions. Again, constitutive rules could be traded in for definitions: the conversational-score function could be defined as that function from conversational-stages to  $n$ -tuples of suitable entities that evolves in the specified way.<sup>33</sup>

As Lewis applies this idea, the elements of the conversational score are things such as sets of presupposed propositions and boundaries between permissible and impermissible actions. Then the acceptability of uttering a particular sentence at a given stage can depend on what is being presupposed. Similarly, the saliencies established by the current score can determine the extension or even the intension of terms such as 'the pig'.

## 2. Deontic Scores and the Pragmatic Significance of Speech Acts

This idiom can be adapted to the model of linguistic practice introduced in this chapter by specifying scores in terms of *deontic statuses*. Linguistic practice as here described can be explained in terms of a score function that determines how the *deontic score* at each stage in a conversation constrains both what performances are appropriate and what the consequences of various performances are—that is, the way they alter the score. The concept of the *pragmatic significance* of a speech act is central to the theoretical metalanguage being employed here. It is a generalization of Dummett's idea of specifying the use of an expression in terms of the pair of its *circumstances* of application and *consequences* of application. (Recall that Dummett's idea was adopted in Chapter 2 as a way of connecting a normative

pragmatics with an inferential semantics.) In scorekeeping terms, the significance of a speech act consists in the way it interacts with the deontic score: how the current score affects the propriety of performing the speech act in question, and how performing that speech act in turn affects the score. Deontic scores consist in constellations of commitments and entitlements on the part of various interlocutors. So understanding or grasping the significance of a speech act requires being able to tell in terms of such scores when it would be appropriate (circumstances of application) and how it would transform the score characterizing the stage at which it is performed into the score obtaining at the next stage of the conversation of which it is a part (consequences of application). For at any stage, what one is permitted or obliged to do depends on the score, as do the consequences that doing has for the score. Being rational—understanding, knowing how in the sense of being able to play the game of giving and asking for reasons—is mastering in practice the evolution of the score. Talking and thinking is keeping score in this sort of game.

In baseball the components of the score (for instance the status a performance can have as a *ball* or a *strike*, or an *out*) are defined in formal terms by the role they play in the process of keeping score—that is, their function in determining what counts as correct play according to the kinematics of score, as codified in the score function. So it is with the components of deontic score in terms of which linguistic practice is to be understood. The deontic statuses of commitment and entitlement are defined in formal terms by the way they can be used to keep track of the moves made in the game of giving and asking for reasons—that is, their function in determining what counts as correct play according to the kinematics of score, as codified in the score function. In this way the notion of *commitment* in linguistic practice plays a role like that of *strike* in baseball: each is an artificial, scorekeeping device.

Besides this formal characterization, however, there is also a material aspect to each of the components in the score, in virtue of which a particular performance qualifies as a ball, a strike, or an out. This material aspect is represented in Lewis's formulation of the kinematics of score quoted above by "the manner *m*," which characterizes the behavior that changes the score from one stage to the next. In games such as baseball, which are not purely formal games (by contrast to chess or tic-tac-toe), the manner in which the score is changed cannot itself be specified entirely in terms of the concepts by means of which the score itself is specified. The complex manner in which a concrete performance qualifies as having the status of a strike or an out invokes such further concepts as the swinging of a bat, the passage of the baseball through a certain region of space specified relative to the position of the batter's body, catching the baseball on the fly, and so on. These further concepts give a material content to the scorekeeping concepts, beyond the formal content they have in virtue of their role in scorekeeping. So it is as well with deontic scorekeeping in linguistic practice. In that case the mate-

rial element concerns such issues as which utterances count as undertaking which commitments, or as deferring to the authority of another asserter, or as invoking noninferential responsive authority.

In baseball the application of scorekeeping vocabulary to particular performances is governed by rules, which are expressed largely in nonscorekeeping vocabulary. The use of this nonscorekeeping vocabulary accordingly answers to norms implicit in the practices of using that vocabulary in contexts other than baseball, where terms like 'inch', 'touch', 'between', and so on already have well-established circumstances of application. Final authority over the application of these rules is vested in the practice of officials. The exact character of this authority is a somewhat complicated matter. There is sometimes an inclination to think of it as *constitutive* of the correct application of the scorekeeping vocabulary, as it is taken to be for instance in the escalating claims of the competing umpires in the familiar tale:

*First Umpire:* I calls 'em as I sees 'em.

*Second Umpire:* I calls 'em as they is.

*Third Umpire:* Until I calls 'em, they ain't.

On such a view, the rules function as something like guides or advisory maxims for the judgment of the umpire, who *makes* a throw into a strike when he *takes* it *as* a strike. But though the attitude of the umpire does determine the status of a throw as a strike for official scorekeeping purposes (that is, does determine what the score is), the use of nonscorekeeping vocabulary in stating the rules that determine how the scorekeeping vocabulary *ought* to be applied to particular cases establishes a perspective from which the judgment of an umpire can nonetheless be understood to be *mistaken*. Metarules explicitly envisage the possibility of such mistakes and, without obliterating their status as such, set up a default-and-challenge system that leaves the umpire entitled to scorekeeping judgments even in the case where they are in fact mistaken, so long as they survive any appropriate challenges that are actually offered.

Linguistic scorekeeping on assertional commitments and entitlements has analogs to both of these dimensions of authority concerning the score. On the one hand, the actual attitudes of scorekeepers are essential in determining the score. On the other hand, the formation of those attitudes is itself subject to norms; scorekeeping is something that can be done correctly or incorrectly. This is not, of course, because it is in general governed by explicit rules; the regress that Wittgenstein and Sellars point to shows that. It is of the utmost importance to the present project to offer an account of what one is doing in taking a scorekeeper to have gotten things wrong, to have attributed commitments different from what the one to whom they are attributed is *really* committed to. For it is in terms of this practical attitude that the possibility of understanding the application of concepts as subject to objective representational constraint—as subject to assessment for being correct



or incorrect in a sense that involves answering to how the things the concepts are applied to actually are, rather than to anyone's attitudes toward them—is eventually to be explained.<sup>34</sup>

One fundamental difference between a game such as baseball and the game of giving and asking for reasons is the *perspectival* nature of the scorekeeping involved in the latter case. As Lewis sets things up, each stage of a baseball game has a single score. One might instead associate a different score with each of the two teams—though some elements, such as the specification of the inning, would be common to both teams at each stage. Linguistic scorekeeping as here construed is more like that: each interlocutor is assigned a different score. For to each, at each stage, different commitments and different entitlements are assigned. There may be large areas of overlap, since almost everyone is committed and entitled to such claims as that  $2 + 2 = 4$ , that red is a color, and that there have been black dogs. But there will also be large areas of difference, if for no other reason than that everyone has noninferentially acquired commitments and entitlements corresponding to different observable situations. These differences ramify both because of the inferential consequences of such observations and because their public availability for inheritance of attitudes and attribution based on testimony varies with conversational exposure. As a result, no two individuals have exactly the same beliefs or acknowledge exactly the same commitments. As with baseball, instead of thinking of these as different scores associated with different interlocutors (corresponding to different teams), they can be aggregated into one grand score for each stage of the conversation of a linguistic community, so long as it is kept clear (in a way corresponding to Lewis's use of subscripts) which deontic statuses are being attributed to which interlocutors.

But linguistic scorekeeping is also perspectival in a way that has no analog at all in baseball. Not only are scores kept *for* each interlocutor, scores are also kept *by* each interlocutor. In baseball there is just one official score, whether it is thought of as the score of the whole game or as the set of scores of each of the teams (and this is, as was pointed out, compatible with there nonetheless being a sense in which the umpire who determines the official score may make a mistake in calling a certain performance a strike). But part of playing the game of giving and asking for reasons is keeping track of the commitments and entitlements of the other players, by *attributing* commitments and entitlements. Just as each interlocutor is typically at each stage attributed a different set of deontic statuses, so each interlocutor typically has at each stage a different set of attitudes or attributions. What *C* is committed to according to *A* may be quite different, not only from what *D* is committed to according to *A*, but also from what *C* is committed to according to *B*. Linguistic scorekeeping practice is *doubly* perspectival.

The idea is that the deontic attitudes of each interlocutor *A* constitute one perspective on the deontic statuses of the whole community. There are, to

begin with, the commitments that *A* acknowledges and the entitlements that *A* claims. Then for each other interlocutor, there are the commitments and entitlements that *A* attributes to that individual. The different sorts of speech acts are to be understood in terms of the different consequences they have for the score that each interlocutor keeps, that is, in terms of how they affect the deontic attitudes of various interlocutors. If *B* asserts that *p*, *B* thereby acknowledges (and so undertakes) a commitment to *p*. So such a commitment ought to be attributed to *B* by anyone in a position to overhear or otherwise find out about that remark.

The pragmatic significance of an assertion goes far beyond this simple shift in deontic attitude on the part of other scorekeepers, however. For the speech act *B* performs has an inferentially articulated *content*, which relates it to other contents. Undertaking commitment to *p* is undertaking commitment as well to its inferential consequences—to those claims *q* that are related to it as conclusions of commitment-preserving inferences having *p* as premise. So if, as a result of *B*'s assertion, *A*'s deontic attitudes change in that *A* comes to attribute to *B* a commitment to *p*, then *A* is obliged also to attribute to *B* commitment to *q*. Or rather, *A*'s treating this as a good inference consists in *A*'s being disposed to keep score in this way, linking the attribution of commitment to *q* consequentially to attribution of commitment to *p*. Again, for *r* to be incompatible with *p* is for commitment to *p* to preclude entitlement to *r*. For *A* to treat these contents as incompatible is for *A* to be disposed to withhold attribution of entitlement to *r* whenever *A* attributes commitment to *p*.

Besides these intercontent, intrapersonal scorekeeping consequences of *B*'s speech act, the assertion may have intraccontent interpersonal consequences regarding *A*'s attitudes. For if *A* takes *B* to be entitled to the claim that *p* (either noninferentially or as the conclusion of an inference), then this may result in *A*'s taking *C* (who also overheard the remark) to be entitled to that claim—but on the basis of testimony, to be defended by deference, rather than either noninferentially or inferentially. The effects of a speech act on the practical attitudes by means of which *A* keeps score on the deontic statuses of various interlocutors depends both on the antecedent score—what they were already taken to be committed and entitled to—and on the content expressed.

### 3. *Inferentially Articulated Significance: Force and Content*

Specifying the pragmatic significance of a speech act kind such as assertion requires showing how the transformation of the score from one conversational stage to the next effected by such a speech act systematically depends on the semantic content of the commitment undertaken thereby. Starting with a notion of the pragmatic significance of speech acts—understood in terms of transformations of the deontic attitudes by which inter-

locutors keep track of each other's commitments and entitlements—it is possible to understand both what it is for two commitments to have the same content and what it is for two commitments to be undertaken by or attributed to the same interlocutor. Not only can these scorekeeping attitudes and shifts of attitude be used to define both contents and interlocutors, the justificatory and communicational links between them can be used to define the notion of *representation*. This is the burden of the discussion of the hybrid deontic attitudes that are made assertionally explicit in the form of *de re* ascriptions of propositional attitudes, in Chapter 8. So the notion of linguistic scorekeeping is intended to play a more fundamental explanatory role here than Lewis has in mind for it. For he is happy to think of conversational scores as kept track of in “mental scoreboards,” consisting of attitudes he calls “mental representations” of the score (representations, presumably, whose content is *that* some component of the score is currently such and such).<sup>35</sup> Clearly he does not envisage a project such as the present one, in which both the nature of mental states such as belief and their representational contents are themselves to be understood in terms of their role in scorekeeping practices, rather than the other way around.

Consider first the notion of the *content* of a speech act or an intentional state. It is motivated first by the idea that speech acts, attitudes, and states of different kinds might share a content—in Fregean terminology, that different sorts of *force* can attach to the same sense. It requires further that the significance of a speech act depends in a systematic way on the content and the sort of force that is attached to it. Dummett's way of putting the point is this: “The implicit assumption underlying the idea that there is some one key concept in terms of which we can give a general characterization of the meaning of a sentence is that there must be some uniform pattern of derivation of all the other features of the use of an arbitrary sentence, given its meaning as characterized in terms of the key concept. It is precisely to subserve such a schema of derivation that the distinction between sense and force was introduced: corresponding to each different kind of force will be a different uniform pattern of derivation of the use of a sentence from its sense.”<sup>36</sup> It does not simply go without saying that such a notion of content is to be had. Use of the theoretical concept of content involves a commitment to displaying the “uniform pattern of derivation of the use from the content,” which Dummett talks about. As he goes on to indicate, one way of reading some of Wittgenstein's remarks is as “rejecting the whole idea that there is any one key idea in the theory of meaning: the meaning of each sentence is to be explained by a direct characterization of all the different features of its use; there is no uniform means of deriving all the other features from any one of them.”<sup>37</sup> This is the point at which semantic theory and pragmatic theory must mesh. It is possible to associate many sorts of things with sentences and other linguistic expressions. What makes the association a *semantic* one is precisely the possibility of appealing to it to explain the

proprieties that govern the *use* of those expressions. Calling what one associates with expressions 'contents', 'propositions', 'sets of possible worlds', 'truth conditions', 'extensions', or 'referents' is at best issuing a promissory note that hints at how what are put forward as their semantic correlates ought to be taken to be relevant to determining how those expressions are correctly used. In the absence of a pragmatics offering an account of what it is to express a content or proposition, to take the actual world to be contained in a set of possible worlds, to try to utter truths, or to employ expressions with various extensions and referents, the theorist's entitlement to the commitment undertaken in treating these associations as semantic is liable to challenge. Semantics answers to pragmatics, attributions of content to explanations of use.

In the present case the pragmatics (comprising the practical proprieties governing linguistic expressions and intentional states alike) is couched in terms of social deontic scorekeeping. The *force* of an utterance, the *significance* of a speech act, is to be understood in terms of the difference it makes to what commitments and entitlements are attributed and undertaken by various interlocutors—that is, in terms of the alteration of deontic scorekeeping attitudes it underwrites. Indeed deontic statuses are to be understood just as ways of keeping such scores. The paradigmatic speech act kind of asserting is specified as having the significance of an undertaking of a commitment (and so the licensing of attributions of that commitment), the licensing or authorizing of further undertakings of such commitments, and the undertaking of a conditional task-responsibility to demonstrate entitlement to the commitment undertaken, if appropriately challenged. This is the sort of significance that must be determined "according to a uniform pattern" (in Dummett's phrase) by the sort of semantic content that is associated with the expressions that qualify as sentences in virtue of their freestanding utterances having this kind of assertional significance.

To be entitled to an *inferential* conception of the contents that qualify as propositional in virtue of their being assertible, then, requires showing how particular assertional significances result from the general account of the speech act of asserting when particular inferential roles are associated with what is asserted. The model of asserting has been constructed with just this criterion of adequacy in mind. To specify the inferential content associated with a sentence, one must, to begin with,<sup>38</sup> indicate the role it plays (in relation to the contents expressed by other sentences) in three different sorts of broadly inferential structure: committive inferences, permissive inferences, and incompatibilities. Doing so is saying what it follows from, what follows from it, and what it precludes or rules out. These are characterized as "broadly" inferential because all of them involve alterations of deontic status that have other alterations of deontic status as their consequences.

In the same sense, the sort of authority that observation reports exhibit counts as broadly inferential because of the reliability inference it involves on the part of the attributor of such authority (discussed in the next chapter).

Although it sounds paradoxical, for this reason the role of a sentence in noninferential reporting should also be understood as falling under the rubric “(broadly) inferential role.” Two features of such specifications are worth focusing on in this context. First, the informal explanations of these inferential relations are in terms of precisely the deontic statuses of commitment and entitlement that are kept track of by the scorekeeping attitudes of interlocutors—that is, just the terms in which the pragmatics is couched. Second, the account of assertional significance in general requires nothing more than inferential roles articulated along these three dimensions in order to determine the significance for social deontic scorekeeping of an assertional utterance.

On the first point, specifying the committive-inferential role of a sentence is specifying the commitment-preserving inferences in which that sentence serves as a conclusion, and those in which (along with other auxiliary hypotheses) it plays an essential role as a premise—essential in that if it is omitted, one could be committed to the remaining premises without therefore counting as committed to the conclusion. This sort of inference is the material-inferential genus of which deductively valid logical inferences are a formal species. Similarly, specifying the permissive-inferential role of a sentence is specifying the entitlement-preserving inferences in which (along with other auxiliary hypotheses) it plays an essential role—essential in that if it is omitted, one could be committed and entitled to the remaining premises (and to no incompatible defeasors) without therefore counting as being entitled to the conclusion. This sort of inference is the material-inferential genus of which inductively good inferences are a species. Incompatibility relations are thought of as broadly inferential because their construal in terms of the deontic statuses of commitment and entitlement is analogous to the first two. The difference is that where the deontic statuses relevant to the two species of strictly inferential relation are homogeneous—both premises and conclusions being assessed in terms of commitments, or both in terms of entitlements—in the case of incompatibility, they are heterogeneous. For to say that two claims are incompatible is to say that *if* one is committed to the first, *then* one is not entitled to the second.<sup>39</sup>

To anticipate the discussion of the next chapter: the sort of default entitlement characteristic of observation reports (perhaps the most important species of this genus) is thought of as broadly inferential because the one who attributes such authority implicitly endorses the reliability of the reporter (under these circumstances and with regard to such contents). Treating someone as a reliable reporter is taking the reporter's *commitment* (to this content under these circumstances) to be sufficient for the reporter's *entitlement* to that commitment. This is endorsing an inference, in the broad sense that corresponds to the consequential relationship between attributing commitment and attributing entitlement to it.<sup>40</sup> Like that involved in incompatibility relations among contents, and in contrast to the two strictly inferential components of content, this inference is heterogeneous with respect to the

deontic statuses involved. Unlike incompatibility, the content of the commitment and the content of the entitlement involved in the reliability inference are the same. The essential role played by endorsement by the attributor of this reliability inference qualifies the role of utterances as elicited observationally as the exercise of reliable differential responsive dispositions for inclusion as contributing to the broadly inferential role of sentences.

The connection between these four kinds of broadly inferential proprieties and the deontic statuses of commitment and entitlement has a double significance. On the one hand, it means that it is possible to understand how social practices of keeping score on commitments and entitlements could confer inferential roles articulated along these four dimensions on the expressions that are caught up in them. That is, social deontic scorekeeping provides an explanation of how expressions must be *used* in order to have contents of this sort associated with them—associated not by the theorist's stipulation but by the practical attitudes of the practitioners whose linguistic conduct is being interpreted. On the other hand, the connection between broadly inferential proprieties of these four kinds and the two sorts of deontic status they involve (homogeneously or heterogeneously) means that once contents articulated in terms of these proprieties have been associated with expressions, it is possible to derive "according to a uniform pattern" the significance that uttering them assertively has for scorekeeping in terms of commitments and entitlements. *Inferring* is accordingly the key concept linking semantic content and pragmatic significance. For not only can propositional semantic contents be understood as inferential roles, but proprieties of inference can be made sense of pragmatically, and specifically assertional significance can be understood in terms of them.

#### 4. *How Inferential and Incompatibility Relations among Contents Affect the Score*

The significance of an assertion of *p* can be thought of as a mapping that associates with one social deontic score—characterizing the stage before that speech act is performed, according to some scorekeeper—the set of scores for the conversational stage that results from the assertion, according to the same scorekeeper.<sup>41</sup> Suppose *A* is such a scorekeeper, and *B* is such an asserter. The way *A*'s score ought to be transformed is settled by the content of *B*'s claim, thought of as its tripartite inferential role in commitment- and entitlement-preserving consequence relations and in incompatibility relations connecting commitments and entitlements. To begin with, *A* must add *p* to the list of commitments attributed to *B* (supposing the more interesting case in which *A* does not already attribute to *B* a commitment with that content). *A* should then add also commitment to any claims *q* that are committive-inferential consequences of *p*, in the context of the other claims attributed to *B*. These will vary, depending on the auxiliary hypotheses available, according to what other commitments *A* already attributes to

*B*. This is closing *A*'s attributions to *B* under commitment-preserving inferences. This closure is determined, in the context of the prior score, by whatever committive-inferential role *A* associates with *p* as part of its content.<sup>42</sup>

Next, the incompatibility relations that *p* (and so its commitment-inferential consequences) stand in must be consulted to determine which, if any, of the entitlements *A* previously attributed to *B* are precluded by the newly attributed commitment. Assertions add new commitments, but they can not only add but also subtract entitlements. Then, in the light of the incompatibility relations associated with all of the commitments attributed to *B*, *A* can attribute *entitlements* to any claims that are *committive*-inferential consequences of commitments to which *B* is already taken to be entitled, closing the attributed score under commitment-preserving inferences, where the resulting attributions of entitlement are not defeated by incompatibilities.<sup>43</sup>

Next, constrained by the entitlement-precluding incompatibility relations associated with all of the other commitments attributed to *B*, *A* can attribute entitlements to any claims that are permissive-inferential consequences of commitments to which *B* is already taken to be entitled, closing the attributed score under entitlement-preserving inferences not defeated by those incompatibilities.<sup>44</sup> Then *A* needs to assess *B*'s entitlement to the claim that *p*, by looking at good inferences having it as a conclusion and premises to which *B* is committed and entitled. This is determined by the first two elements of the broadly inferential role *A* associates with *p*. Similarly, *A* must assess the possibility of *B*'s noninferential default entitlement to *p*. Whether *B*'s undertaking of commitment to *p* falls within the scope of any good reliability inferences, according to *A*, depends on what else *A* is committed to—concerning the conditions under which the deontic status was acquired, implicitly, whether they are among those (if any) in which *B* is a reliable reporter with respect to contents such as *p*. Again, *A* must assess *B*'s entitlement to *p* as testimony, by inheritance of the entitlement *A* attributes to some other interlocutor (possibly even *A*) who has asserted it at an earlier stage. If *A* takes *B* to be entitled to *p* by any of these mechanisms of inheritance and default, then *A* will take *B* to have successfully entitled others (including *A*) to that claim (in the absence of incompatible defeasors). In this way the broadly inferential content that *A* associates with *B*'s claim determines the significance *B*'s assertional speech act has from the point of view of *A*'s scorekeeping, the difference it makes to the deontic attitudes of attributing and acknowledging commitments and entitlements by means of which *A* keeps track of everyone's deontic statuses.

##### 5. *Deferrals, Disavowals, Queries, and Challenges*

The model of assertion defined by a scorekeeping function that appeals in this way to broadly inferential assertible contents to determine the significance of assertions of those contents can be enriched by allowing

various auxiliary sorts of speech acts. *Deferrals* have already been mentioned. No new sorts of content need be considered in order to specify the significance of deferring for social deontic scorekeeping. The content associated with a deferral is just the assertible content of the commitment that the deferrer is seeking to vindicate by indicating a testimonial path whereby entitlement to it can be inherited. It is the force associated with that content that is different from the assertional case.

*A* assesses *C*'s deferral to *B* concerning *p* by assessing first *B*'s entitlement to *p* (as considered above), and then *C*'s entitlement to inherit it. This latter is a matter not only of its being the case that *A* does not attribute to *C* commitment to anything incompatible with *p* (an issue in general independent of whether *A* attributes to *B* commitment to anything incompatible with *p*). For even if *A* does not attribute to *C* commitments incompatible with the *claim B* made (the commitment undertaken), it is still possible that *A* attributes to *C* commitments incompatible with inheriting *B*'s entitlement to it. This would happen if *B*'s entitlement, according to *A*, depends on justifying the claim *p* by appeal to the claim *q*, where *A* takes it that *C*, but not *B*, is committed to some claim incompatible with *q*. Similarly, it might be that *C* is committed, though *A* is not, to some claim incompatible with one of the conditions (according to *A*) for *B*'s observational authority with respect to *p*. Thus if *C* takes it that *B* is looking through a tinted window, *A* may take this to preclude *C*'s inheritance of entitlement to *B*'s noninferential report of the color of a piece of cloth, even though *A* takes it that *C* is wrong about the conditions of observation.

*Disavowals*, *queries*, and *challenges* are three other speech acts auxiliary to assertion that it is useful—from a scorekeeping point of view—to include in a model of the game of giving and asking for reasons. Disavowals permit one to repudiate or disclaim a commitment one has previously undertaken or to make it clear that one does not acknowledge such a commitment. Again no new sorts of content need to be considered; the force or significance of speech acts of this sort (the difference they make to the score that interlocutors keep on each other's deontic statuses) is determined by associating with them broadly inferential, assertible contents of the sort already discussed. It is the force of the speech act to which such contents attach that is different. For *A* to take *B*'s disavowal of commitment to *p* to be successful is for *A* to cease attributing commitment to *p* to *B*, and to reinstate any attributed entitlements that were withheld because they were defeated by their incompatibility with *p*.

Why might a disavowal not be successful? Because of the two fundamental ways *B* can undertake a commitment (and thereby license *A* to attribute it)—about which more below. For *B* might acquire commitment to *p* directly, by avowing it—that is, by overtly asserting it. Or *B* might acquire that commitment indirectly, as a consequence of a commitment (perhaps itself avowed) to *q*, from which it follows by a good commitment-preserving infer-



ence (according to *A*). In such cases, *B*'s disavowal of *p* can be successful (according to *A*) only if *B* is also prepared to disavow *q*. Indeed, disavowing *p* is indirectly disavowing *q*. But if *B* persists in asserting *q*, that commitment is incompatible with the disavowal, and the disavowal of *p* cannot accordingly reinstate entitlement to claims *A* withholds attribution of entitlement to only because of the flaw in *B*'s title represented by commitment to the incompatible *p*. Again, the sort of significance for *A*'s scorekeeping that disavowals of *p* have requires no further elements of content beyond those involved in its assertional use, even though the significances of the two sorts of speech act are quite different—indeed, in some ways complementary.

It would also be useful to those keeping score if there were some way of eliciting the avowal or disavowal of a particular claim—a way for *A* to find out whether *B* acknowledges commitment to *p*. Such a speech act is a basic query: *p*? By itself, such a speech act would have no effect on the deontic score; only responses to speech acts of this kind would alter the score. In the basic model, there is no reason not to allow anyone to be entitled to such a query at any point in a conversation.

Another sort of speech act that might be distinguished is *challenging* the testimonial authority of an assertion. As was indicated above, this might consist in no more than making an incompatible assertion. But it might be useful from a scorekeeping point of view to have a way of addressing an assertion as a challenge to another assertion. The significance of such a challenge is to bring attributions of entitlement by default into question wherever the challenging assertion is one the challenger is at least *prima facie* entitled to. For *A* to treat *C*'s challenge of *B*'s assertion of *p* as successful is for *A* to respond to it by withholding attribution of entitlement to *B* for that claim, pending *B*'s vindication of it, whether inferentially or deferentially. This has the effect of making that assertion unavailable (according to *A*'s score) to other interlocutors who might otherwise inherit entitlement to commitments to the same content testimonially from *B*. There is no reason in principle that conflicts of this sort need to be resolvable. The public status of competing claims may remain equivocal in that neither the challenged nor the challenging claim can be vindicated successfully, or in that both can be—though of course *A* will not take it that any one interlocutor could inherit entitlements to commitments to both of the incompatible contents. "Let a thousand flowers blossom. Let a hundred schools of thought contend."

## 6. *Acknowledged and Consequential Commitments*

In the next chapter the model of assertion is enriched by adding another variety of discursive commitment. Besides the cognitive commitments undertaken by assertions, *practical* discursive commitments—that is, commitments to *act*—are considered. The speech acts that undertake such commitments, namely expressions of intention, have quite a different sort of

significance from assertions. Attributions of them (the attitudes in terms of which score is kept on these deontic statuses) also behave differently. Yet the significances of these further sorts of performances, statuses, and attitudes can be understood straightforwardly by analogy to the sorts of scorekeeping that have been introduced for the pure assertional case. That extension of the model of the game of giving and asking for reasons provides a way of understanding intentional action and intentional interpretation of agents. It has been suggested that the doxastic commitments undertaken by speech acts having the significance of assertions can serve as analogs of belief—that such deontic statuses can do much of the explanatory theoretical work usually done by the paradigmatic sort of intentional state. That claim clearly cannot be assessed until the model is extended so as to include the possibility of manifesting such commitments not only in what is *said* but also in what is *done*—in action as well as assertion. Before turning to that wider context, however, it is worth pausing briefly to consider what can be said about the relation between belief as an intentional state and the deontic statuses with assertible contents as considered so far, already in the more restricted context of purely assertional practice.

It was indicated above that the assertional practices described so far generate two different senses in which one could be taken to be committed to a claim. Interlocutors undertake some commitments directly, by avowing them overtly: performing speech acts that have the significance of assertions. The commitments one is disposed to avow<sup>45</sup> are *acknowledged* commitments. But in virtue of their inferentially articulated conceptual contents, assertional commitments have consequences. Undertaking a commitment to a claim with one content involves undertaking commitments to claims whose contents are (in the context of one's other commitments) its committive-inferential consequences. Undertaking commitment to the claim that Pittsburgh is to the West of Philadelphia is one way of undertaking commitment to the claim that Philadelphia is to the East of Pittsburgh. These *consequential* commitments may not be acknowledged; we do not always acknowledge commitment to all the consequences of the commitments we do acknowledge. They are commitments nonetheless. For the only way that deontic *statuses* enter into the scorekeeping specification of assertional practices is as the objects of deontic *attitudes*. Indeed, *all* one can do with a commitment (or entitlement), in the model presented here, is take up a deontic attitude toward it—attribute it or undertake it, either directly by acknowledging it, or indirectly and consequentially. The scorekeeping model trades in talk about the status of *being* committed for talk about proprieties of practical attitudes of *taking* to be committed. Deontic statuses are just something to keep score with, as balls and strikes are just statuses that performances can be treated as having for scorekeeping purposes. To understand them, one must look at actual practices of keeping score, that is, at deontic attitudes and changes of attitude.

These ways in which one can come to be committed to a claim—by acknowledgment and consequentially—correspond to two ways in which we talk about *belief*. In one sense, one believes just what one takes oneself to believe, what one is prepared to avow or assert. In another sense, one believes, willy-nilly, the consequences of one's beliefs. Believing that Pittsburgh is to the West of Philadelphia is believing that Philadelphia is to the East of Pittsburgh, whether one knows it or not. This second sense looms particularly large for those who take their role in intentional explanations of behavior as the touchstone for identifying beliefs. For such explanations work only to the extent the individual in question is (taken to be) *rational*—not to have contradictory or incompatible beliefs and to believe the consequences of one's beliefs. From the point of view of the present project, this is because the conclusion of intentional explanations, strictly construed, is always a *normative* one—given intentional states whose contents are thus and so, one *ought* rationally or is rationally *obliged* or *committed* to act in such and such a way. Drawing conclusions about what actually will be or was done requires an additional premise, to the effect that the individual in question has mastered the practices of giving and asking for reasons sufficiently to be disposed to respond to acknowledgment of a (practical) commitment by producing a performance that satisfies it.

The sense of belief in which one is taken actually to believe what one ideally *ought* to believe (at least given what else one believes), call it *ideal* or *rational* belief, can conflict with the sense of belief for which avowal is authoritative. Dennett distinguishes these but thinks of them as competing norms to which a univocal sense of 'belief' must answer: "These two interdependent norms of belief, one favoring the truth and rationality of belief, the other favoring accuracy of avowal, normally complement each other, but on occasion can give rise to conflict."<sup>46</sup> The conflict arises precisely because one can avow incompatible beliefs, and fail to avow even obvious consequences of one's avowals:

What better source could here be of a system's beliefs than its avowals? Conflict arises, however, whenever a person falls short of perfect rationality and avows beliefs that are either strongly disconfirmed by the available empirical evidence or are self-contradictory or contradict other avowals he has made. If we lean on the myth that man is perfectly rational, we must find his avowals less than authoritative: "You *can't* mean—understand—what you are saying!"; if we lean on his 'right' as a speaking intentional system to have his word accepted, we grant him an irrational set of beliefs. Neither position provides a stable resting place; for, as we saw earlier, intentional explanation and prediction cannot be accommodated either to breakdown or to less than optimal design, so there is no coherent intentional description of such an impasse.<sup>47</sup>

The notion of incompatible beliefs offers no difficulties for a normative construal of intentional states as deontic statuses. There is nothing incoherent or unintelligible about the idea of undertaking incompatible commitments—incompatible or inconsistent beliefs just go into a box with incompatible or inconsistent promises. This is one of the benefits of this sort of approach over causal-functional accounts of intentional states. Yet the tension that Dennett identifies is a real one. The decision to treat belief just as what one is prepared to avow “amounts to the decision to lean on the accuracy-of-avowal norm at the expense of the rationality norm . . . If we demand perfect rationality, we have simply flown to the other norm at the expense of the norm of accuracy of avowal.”<sup>48</sup>

Dennett does not offer any way to reconcile the competing demands that the norm of rationality and the authority of avowals place on attributions of belief. The terminology employed here is animated in part by the thought that ‘belief’ may simply be *ambiguous* between a sense in which one believes just what one is prepared to avow and a sense in which one also believes what one *ought* rationally to believe, as a consequence of what one is prepared to avow (as already indicated, failures of rationality due to incompatibility cause no particular trouble once intentional states are construed as deontic statuses). An unambiguous, univocal technical term ‘doxastic *commitment*’ is introduced, which comprises both commitments one is prepared to avow and commitments that follow from those one acknowledges. But attention to the attitudes in terms of which those deontic statuses are explained makes it possible also to distinguish clearly between these two *kinds* of commitment, as ‘belief’-talk does not. The proposal is accordingly not to *analyze* belief in terms of commitment but to discard that concept as insufficiently precise and replace it with clearer talk about different sorts of commitment.

The fundamental concept of the metalanguage employed in specifying the model of assertional practice is that of the deontic attitude *attributing* a commitment. For the deontic attitude of *undertaking* a commitment is definable in terms of attribution: undertaking a commitment is doing something that licenses or entitles others to attribute it. Assertional performances or avowals are performances that express the deontic attitude of acknowledging doxastic commitments. They license attribution of (and insofar as they are successful, deferral with respect to) both the commitments they express and those whose contents are appropriate inferential consequences of the contents, commitment to which is overtly acknowledged. The attitude of acknowledging a commitment is in effect that of attributing it to oneself.<sup>49</sup> The fact that one thereby undertakes consequential commitments that may reach beyond what one acknowledges just shows that the generic attitude of *undertaking* a commitment is not to be identified with its species *attributing* a commitment to oneself, which is acknowledging it. The *social* dimension (invoking the perspective of *other* attributors) is essential to understanding

undertaking in terms of proprieties of attributing. The way in which the collaboration of attitudes adopted from two socially distinct perspectives—attributions of commitment to oneself and by others—is required to institute discursive commitments is the central theme of this work. It is in terms of the social-perspectival character of discursive deontic statuses that the notion of *objectivity* is to be made intelligible—both the general normative distinction between what one is *really* committed to do (or ought to do) and what one is merely *taken* by someone to be committed to do, and the more specific version that underwrites the notion of objective representational content, of a claim's correctness answering to how things are with what it represents, rather than to what anyone *takes* to be correct.

The roots of this social-perspectival account can already be discerned in the distinction that a scorekeeper can make between the commitments an interlocutor has undertaken and those that interlocutor acknowledges, and so is prepared to assert. For the attributions of the scorekeeper distinguish between the actual deontic *status* of the one for whom score is being kept, what that interlocutor is *really* (consequentially) committed to, and the deontic *attitudes* of that subject, what that interlocutor acknowledges commitment to by self-attribution. In other words, the notion of consequentially undertaking commitments provides the basis for distinguishing (in terms of the attitudes of someone keeping score) between deontic statuses and deontic attitudes. Indeed (as will appear in Chapter 8, where this issue is explored) the very notion of one propositional content being an inferential *consequence* of another essentially involves a crucial relativity to social perspective: are the auxiliary hypotheses (the premises to be conjoined with the claim in question in assessing its consequences) to be those the scorekeeper assessing the propriety of the inference *undertakes* commitment to, or those the scorekeeper *attributes* to the one whose statuses are being assessed?

Neither answer is correct. The fact that proprieties of inference a claim is involved in can be assessed from either of two social perspectives—that of the one *attributing* commitment to the claim or that of the one *undertaking* that commitment—is fundamental to the very notion of a *propriety* of inference. And since propositional and so conceptual contents of all sorts are constituted by the broadly inferential proprieties of practice in which they are caught up, such contents are essentially social and perspectival in nature. The propositional content of a claim or commitment can be specified only from some point of view; that it would be differently specified in definite ways from other particular possible social perspectives (that is, scorekeepers occupying such perspectives) is an essential part of its being the content it is.

At this point the phrase “social-perspectival character of the contents of discursive commitments” can be little more than a label attached to a promissory note—though the discussion of scorekeeping in this section is intended to give it enough resonances to make it at least a suggestive label. Before that promissory note is redeemed, in Chapter 8, it is necessary to look

much more closely at the sorts of contents that can be conferred on expressions playing various roles in assertional practices of the sort described here. In the intervening chapters the notion of *inferential* articulation is deepened and extended by adding *substitutional* machinery. The notion of substitution inferences permits the extension of the notion of conceptual content introduced here to essentially *subsential* expressions such as singular terms and predicates, which can play only indirectly inferential roles—not serving themselves as premises and conclusions of inferences but only occurring in the sentences that can serve in those capacities. The conceptual content expressed by the use of singular terms and predicates is articulated by substitution-inferential commitments (in Chapter 6). Extending this account to the sort of content expressed by the token-reflexive or indexical use of unrepeatable expressions, paradigmatically the sort of *deictic* tokenings that play such an important role in empirical knowledge claims, requires looking still further at *anaphoric* connections among tokenings. (The importance of anaphoric relations for understanding what is expressed by traditional semantic vocabulary, paradigmatically ‘true’ and ‘refers’, is argued in Chapter 5.) Anaphora is explained as a certain structure of *inheritance* of substitution-inferential commitments (in Chapter 7). The result is a layered account of the semantic contents that can be conferred by assertional practice as here described—an account whose key concepts are those of *inference*, *substitution*, and *anaphora*, each permitting a finer-grained analysis of the structures that precede it and which it presupposes.

# Perception and Action: The Conferral of Empirical and Practical Conceptual Content

---

Nature to be commanded must be obeyed; and that which in contemplation is as the cause is in operation as the rule.

BACON, *Novum Organum*

In the beginning was the deed.

GOETHE, *Faust*

The true being of man is his deed; in this the individual is actual . . . What the deed is can be said of it. It is this, and its being is not merely a sign, but the fact itself. It is this, and the individual human being is what the deed is. Action simply translates an initially implicit being into a being that is made explicit.

HEGEL, *Phenomenology*

---

## I. ASSERTIONS AS KNOWLEDGE CLAIMS

### 1. *Five Strategic Explanatory Commitments*

Assertions are the sort of *claims* made in the standard case by uttering freestanding declarative sentences—that is, sentences whose occurrence is not embedded in the occurrence of a compound sentence. A commitment has been undertaken here to an order of explanation dictating that this principle be exploited by defining declarative sentences in terms of an account of assertion, which evidently then is required to be made available independently. This contrasts with the procedure common in formal semantics, in which the theorist leaves until later the task of getting a grip on the activity, force, or significance of assertion but provides an antecedently defined construal of sentences. The strategic commitment to treating what is expressed by the use of *sentences* (rather than what is expressed by the use of singular terms or predicates) as the fundamental sort of semantic content is an element of the present account that has been taken over from Kant.

The pragmatist strategic commitment to understanding semantics in terms of pragmatics (the *contents* associated with expressions in terms of the practices governing their *use*) is an element of the present account that has been taken over from Wittgenstein. The strategic commitment to specifying such a pragmatics in the first instance in *normative* terms is an element of

the present account that has been taken over from Kant, Frege, Wittgenstein, and Sellars. The inferentialist strategic commitment to treating the public linguistic practice of *asserting* as the fundamental activity involving such contents, rather than the private mental practice of judgment, is an element of the present account that has been taken over from Dummett. The strategic commitment to understanding asserting a sentence as a significance a performance acquires in virtue of its role in a practice of giving and asking for *reasons*, of justifying and communicating justifications, is an element of the present account that has been taken over from Sellars.

This constellation of commitments combines the *normative* articulation of the pragmatic significances of assertional performances with the *inferential* articulation of the propositional contents they express, in part by putting the issue of whether an asserter is *entitled* to the *commitment* undertaken by making an assertion at the center of the practice that institutes those significances and confers those contents. To do so is to treat the sort of claim involved in asserting as an implicit *knowledge* claim. From the point of view of the concerns that motivate the present project, this is as it should be. For the aim all along has been to elaborate a criterion of demarcation that sets us off by our peculiar susceptibility to *reasons*. It is this susceptibility that makes it appropriate to think of ourselves in terms of the categories of *knowledge* and *action*. That is why the story really begins with Kant's observation that knowings and actions are to be distinguished from other things we do by the characteristic way in which we are *responsible* for them. The notion of *discursive commitment* arises in the domain of social practice when one focuses specifically on the norms that are articulated in the form of reasons.

Absent the inferential dimension, the norms implicit in a set of social practices could be understood neither as conferring *propositional* contents nor as instituting *assertional* significances, hence not as governing genuinely *linguistic* practice. Inferential connections enter into the alterations of attitude (the scorekeeping that defines assertional practice) in three fundamental ways: one corresponding to each of the basic sorts of deontic status, and a third involving the relation between them. First, part of the significance of acknowledging an assertional commitment is that one thereby undertakes commitment as well to all those contents it entails—that is, to consequences that follow from it by commitment-preserving inferences. One who claims that a lion roared is committed thereby to a mammal's having roared. Second, part of the significance of undertaking an assertional commitment is that one thereby undertakes a conditional task-responsibility to demonstrate one's *entitlement* to that commitment, if faced with a warranted challenge. Here justificatory or entitlement-preserving inferences involving the asserted content help determine what deontic statuses are attributed to which asserters, challengers, and deferrers. Finally, part of the deontic scorekeeping practice within which performances can have the significance characteristic of claim-



ings is to withhold attribution of *entitlement* to commitments *incompatible* with a *commitment* that has been undertaken (whether by overt assertion or consequentially). This is the practice that defines incompatibility relations on the contents of deontic states; two claims are incompatible if commitment to one precludes entitlement to the other. Connections of these three sorts consist in proprieties that govern the alterations of deontic attitudes by which interlocutors keep discursive score. Deontic statuses then count as having inferentially articulated contents because of the pragmatic scorekeeping significance of the performances that express the acquisition of those statuses.

## 2. Knowledge as a Complex Hybrid Deontic Status

That assertions have the default status and significance of implicit knowledge claims is to be understood in terms of these inferentially structured interactions between the two *modally* distinct deontic statuses, *commitment* and *entitlement*, and the two *socially* distinct deontic attitudes, *attributing* and *undertaking* deontic statuses. The status one attributes in attributing knowledge is traditionally understood according to the tripartite structure: justified true belief (JTB).<sup>1</sup> One of the leading ideas of the present approach is that the notion of normative status can be made to do much of the theoretical and explanatory work that the notion of intentional state has heretofore been called on to do. In the social practice model, talk of *belief* is replaced by talk of assertional or doxastic *commitment*. According to the JTB approach, attributing knowledge is attributing a special kind of belief. So attributions of knowledge are to be rendered here in terms of the deontic attitude of *attributing commitments*—specifically commitments of the sort that can be undertaken or acknowledged by performing a speech act that has the significance characteristic of assertions.

As Plato had already pointed out, there is more to attributing knowledge than attributing belief. There is also the issue spoken to by the demand for justification, or as Plato has it, for an account. According to the canonical tripartite understanding, knowledge is not just belief but *justified* belief. Clearly what corresponds to this condition in the deontic version is the demand that the one taken to be a knower not only have a commitment but be *entitled* to that commitment. Making an assertion, it has been said, is making a knowledge claim. Assertional performances as modeled here have the significance not only of undertaking commitments but of defeasible claims to entitlement to those commitments. So one is not attributing knowledge to someone unless one not only attributes a commitment but also attributes a corresponding entitlement. Most classical epistemological problems are really problems concerning this deontic status—justification for believing, or more generally entitlement to believe.

Before considering entitlement to believe, however, a few words are in

order about the third limb of the tripartite rendering of what one is doing in taking someone to have the status of a knower. In attributing knowledge one is not just attributing justified belief—a commitment of the sort that can be undertaken by asserting it and an entitlement to that commitment of the sort that can be demonstrated by justifying it. One must also take the belief to be *true*. What is the social-deontic attitude corresponding to the truth condition on attributions of knowledge?

The attitude of taking-true is just that of *acknowledging* an assertional commitment (the attitude that grounds consequential undertakings of such commitments). A theory of asserting and assertional commitment is a theory of taking-true. Evidently this principle can be exploited according to two different orders of explanation: moving from a prior notion of truth to an understanding of asserting (or judging) as taking, treating, or putting forward *as true*, or moving from a notion of asserting to a notion of truth as what one is taking, treating, or putting forward a claim *as*.<sup>2</sup> The latter line of thought accords 'true' an expressive role, in permitting us to *say* something about assertion, rather than an explanatory role, as something that can be understood in advance of understanding assertion and used to advance to such an understanding. This approach is the one pursued in Chapter 5.

In taking someone to be a knower, one *attributes a commitment, attributes entitlement* to that commitment, and *acknowledges commitment* to the same content oneself. Undertaking the commitment is part of what the asserter authorizes others to do—not only to attribute the commitment but also to undertake it, on the asserter's authority. That authority depends on the asserter's entitlement to the commitment, so the asserter is implicitly claiming that entitlement as well. That is why assertions in the basic practices described here have the significance of claims to knowledge. For others to take those claims to be successful is for them to attribute the commitment undertaken, in addition to attribute entitlement to it, and finally to endorse the claim themselves. These correspond, in the model of linguistic practices in terms of scorekeeping with deontic attitudes, to taking to believe, taking to be justified in that belief, and taking the belief to be true.<sup>3</sup>

According to this way of understanding knowledge claims, the significance of the truth condition on attributions of knowledge lies in the fundamental difference in social perspective between *attributing* a commitment (or other deontic status) to another and *acknowledging* it oneself. Knowledge is a *complex deontic status*, in the sense that it involves both commitment and entitlement. But attributions of knowledge (and so claims of knowledge) are also *hybrid deontic attitudes*. So knowledge can be called a hybrid deontic status. Attributions of knowledge (the attitudes in terms of which that status is to be understood) are hybrid deontic attitudes in the sense that they involve both attributing and acknowledging commitments. These attitudes are perfectly intelligible in the context of the model presented here of the social practices that institute assertional force.

But it is also clear in those terms that there is a great danger of misinterpreting what one is doing in calling a claim true. The danger is in misunderstanding taking-true solely in terms of *attributions* of status (rather than as essentially involving also the *undertaking* of one), by assimilating that attitude too closely to the attributions of commitment and entitlement involved in the other dimensions involved in treating a claim as a bit of knowledge. If the hybrid nature of that attitude is overlooked, it will be thought of as consisting only in attributions; attributions of knowledge will be taken as comprising attributions of belief, attributions of justification, and attributions of truth. One then looks for the property or status one is attributing to a claim in taking it as true.<sup>4</sup> The property or status projected by misconstruing undertaking a commitment as attributing some property or status to it is bound to be “queer.”

In the deontic scorekeeping model of inferentially articulated linguistic social practices, asserting is making a knowledge claim. The attitudes in terms of which the hybrid deontic status of knowledge is understood are just those in terms of which the significance of assertions is specified. Assertional practice is accordingly a version not only of the game of giving and asking for reasons but also of the game of making and defending claims to knowledge. Practitioners who can produce and consume assertions are *linguistic* beings. Practitioners who can produce and consume reasons are *rational* beings. Practitioners who can produce and consume knowledge claims are *cognitive* beings. On the account presented here, these are three ways of talking about the same practices and the same capacities.<sup>5</sup>

Underlying all of them is the inferentially and socially articulated notion of *discursive commitment*. It is the topic in which philosophy of language, philosophy of mind,<sup>6</sup> and epistemology are alike rooted. What epistemology studies is a deontic status that is implicitly in play in any practices involving propositional contents—whether or not those practices include the expressive resources provided by words like ‘knowledge’, which can be used to make attitudes toward that status explicit. For making and defending what are implicitly claims to knowledge is an essential feature of discursive practice as such.

On this account, prizing and searching for knowledge are not specialized intellectual virtues, appropriate only to a sophisticated, culturally late-coming elite. They are built into what we fundamentally are. The complex hybrid deontic status of knowledge defines the *success* of assertion. Treating an assertion as expressing knowledge—attributing to the asserter entitlement to the commitment undertaken thereby and endorsing that commitment oneself—is the response that constitutes the practical recognition of the authority implicitly claimed by the assertion. For that is the authority to license undertakings of commitment to that same claim by those in the audience, in virtue of the asserter’s entitlement to the commitment. For a scorekeeper fully to accept the authority implicitly claimed in the making of an assertion

is just for that scorekeeper to treat it as having the status of knowledge. So the aspiration not only to *truth* but to *knowledge* is built right into the normative structure of assertional practice. (And it should be noticed that the sense in which the status of knowledge provides the *point* of assertion can be specified in advance of any consideration of the *intentions* of the practitioners.) Knowledge is on this account an ideal projected by the very possibility of saying anything at all.

### 3. *Justifying and Being Entitled*

In this context it is useful to look a bit more closely at the structure of attributions of entitlement to assertional commitments, from a more traditionally epistemological perspective. According to the tripartite analysis, to take a claim to express knowledge is to take it to express a justified true belief. The justification condition on knowledge will be misunderstood if one does not distinguish between two senses in which a belief can be said to be justified. In one sense, to call a belief justified is to invoke its relation to the process of *justifying* it. To *be* justified in this sense is to *have been* justified—exhibited as the conclusion of an inference of a certain kind. In another sense, to call a belief justified is to attribute to it what might be called *positive justificatory status*. Positive justificatory status is just what has been talked about here in terms of *entitlement* to a claim.

The relation between possession of such status and the activity of justifying may be quite indirect. In particular, justifying a claim is only one way in which it can acquire positive justificatory status. Indeed, as has already been pointed out, to avoid embarking on a foundationalist regress it is necessary to acknowledge that a commitment may have a positive justificatory status without having been justified (indeed, without that entitlement having been defended in any way, whether *intrapersonally* by inference or *interpersonally* by deference). Since any activity of justifying—even if that term is understood broadly, as entitling (so as to include deferring as well as inferring)—is a mechanism making possible the *inheritance* of entitlements, there must be some at least *prima facie* entitlements available to get the process off the ground. If dogmatism is to be avoided, such entitlements must not be immune to criticism; there must be mechanisms for bringing them into question. The combination of *prima facie* entitlements and ways of criticizing and undermining them is what was called the *structure of default and challenge*. It characterizes a dynamic process of acquisition and loss of entitlements by various commitments on the part of various interlocutors (kept track of in the attitudes of claiming and attributing entitlements) and of withholding such claims and attributions.

Classical foundationalism considers only justifying in the narrow sense of an inferential activity, not in the broader sense of vindication that includes the communicational dimension appealed to by deferential entitling (the

authority of testimony). This is unfortunate, for if the analysis just offered of what one is doing in calling something knowledge is correct, the hybrid deontic status of knowledge is incomprehensible in abstraction from the social distinction of perspective distinguishing the deontic attitudes of attributing and undertaking commitments. One of the centerpieces of the present account is its attention to the interaction of the two dimensions of the practice of giving and asking for reasons for commitments to inferentially articulated contents: the intracontent, interpersonal communicational dimension and the intercontent, intrapersonal justificatory dimension. Both the individuation of the contents individuals are responsible for and the individuation of the individuals responsible for them are to be understood in terms of this structure. Equally important, as Chapter 8 shows, an inferential understanding of the *representational* dimension of conceptual content depends upon an appreciation of the *social* articulation of inferential practice. None of this is accessible from the point of view of the one-dimensional approach that ignores the significance of communication for justification. Even within that narrower compass afforded by exclusive attention to intrapersonal, intercontent entitling, however, the consequence of insisting that positive justificatory status can be the result only of justifying is a dual regress—one regress on the side of entitlement to premises, and another on the side of entitlement to inferences.

To illuminate the default and challenge response to the threat of a foundationalist regress on the side of premises, consider its twin on the side of inferences. If entitlement to a commitment to  $q$  is at issue, and that commitment is justified by asserting  $p$ , the vindication might be unsuccessful either because the commitment to  $p$  is not one the interlocutor is entitled to or because the inference from  $p$  to  $q$  is not correct (in this case, not entitlement-preserving). A regress on the side of the inferences results if one insists that each inference is, to begin with, in need of support or justification. Endorsing the propriety of an inference is brought into the game of giving and asking for reasons in a new way by making the inference explicit in the form of a conditional, which can be endorsed, challenged, and defended like any other assertible. The demand is then for an explicit rule or principle to warrant the propriety of every inferential transition appealed to in justifying a claim.

But pragmatism maintains that to demand this is to view things the wrong way around. One must start with a notion of taking or treating inferences as correct in practice. Without such a practice, there is no game of giving and asking for reasons to bring inferences into in the form of explicit assertions. Once the game is under way, the practical inferential attitudes it involves can then, on suitable occasions, be made explicit in the form of endorsements of conditionals. But what those conditionals express is intelligible only in terms of the underlying inferential practice.

If it is insisted instead that no move be treated as entitled or entitlement-

preserving until its entitlement has been demonstrated or justified, a new premise is introduced corresponding to every inference, and also a new inference employing that premise. Then the regress ramifies, as the entitlements to those new premises and those new inferences must themselves be secured. Within the resulting regress can be discerned Kant's and Wittgenstein's regress of rules—and where in addition the goodness of inference is identified with *formal* goodness of inference, Lewis Carroll's regress of conditionals and detachment from conditionals as well. Carroll's multiplication of premises standing behind inferences should be halted by acknowledging primitive *rules* of inference; the multiplication of conditionals explicating implicit "enthymematically suppressed" premises should be halted by acknowledging primitive *material* rules of inference; and the multiplication of rules should be halted by acknowledging primitive material-inferential *practices*.

These moves have all already been considered. The default-and-challenge structure of assertional entitlement just amounts to extending to the case of assertions the policy that underwrites these ways of thinking about inferences. What is fundamental in each case is the practical attitude of taking or treating as correct moves in the game of giving and asking for reasons. Though such entitlements can be brought into question later, one initially is entitled to whatever one is in practice taken or treated as entitled to; deontic statuses must be understood in terms of practical deontic attitudes. It makes sense that this way of construing the proprieties of inference that articulate the propositional content of assertional speech acts and the commitments and entitlements they involve should extend as well to the proprieties that govern those assertional performances and deontic statuses. For asserting and inferring are two sides of one coin; neither activity is intelligible except in relation to the other. Undertaking an assertional commitment involves a commitment to the propriety of inferences from the circumstances of application to the consequences of application of the concepts in terms of which its content is articulated. If claiming is to be possible at all, some of those content-constitutive implicit inferential proprieties must in practice be taken for granted, treated as *prima facie* in order—not as innocent until *proven* guilty, but at least as innocent until indicted on the basis of reasonable suspicion. In the same way, sometimes a defeasible presumption that the application of those concepts in an assertion or judgment is appropriate must be in order.

## II. RELIABILITY

### 1. *Reliabilism and Entitlement*

When the justification-as-entitlement of a belief is decoupled to this extent from the activity of inferentially (or, for that matter, deferentially) justifying it, the question arises whether the latter notion need be taken to

play any role whatever in the understanding of the status of being a justified belief that is appealed to by the tripartite analysis of knowledge. It is generally agreed that *some* sort of entitlement to a claim is required for it to be a candidate for expressing knowledge. But it is not obvious that inferring in the sense of justifying is at all fundamental to that sort of entitlement. When examples of the sort that motivate the tripartite analysis in the first place are examined more closely, it appears that what is forbidden is that it be merely *accidental* that one has a true belief. Someone who flips a coin to decide which is the correct road to Athens may by accident pick the right one, and may somehow or other come to believe that the one chosen is the correct road. But such a true belief does not qualify as knowledge. One way of showing that the belief is not merely accidental is indeed to provide an account, to offer reasons for the belief.

It has been suggested, however, that this is merely *one* way, and by no means the most basic, in which the belief could be shown to have credentials beyond those provided by happenstance and coincidence. In particular, the correctness of the belief is not merely fortuitous if it is the outcome of a generally reliable belief-forming mechanism. Epistemological *reliabilists* claim that this is the sort of entitlement status that must be attributed (besides the status of being a true belief) for attributions of knowledge.<sup>7</sup> The perceptual mechanisms underlying entitlement to empirical claims provide the most important and persuasive examples. This line of thought is sometimes extended to an analysis of justification as consisting simply in the demonstration of the reliability of a belief-forming mechanism. The version that is of interest here, however, claims only that reliability can do *all* the work that inferential justifying is taken to do in the standard tripartite analysis (and, by extension, in the assertional practices of making and defending knowledge claims described here). The further Procrustean reductive assimilation of justifying to the paradigm of indicating reliability can be put to one side at this point, for it is the difficulties that arise already with the weaker thesis that are most instructive.

*Reliability* accounts of entitlement to assertional commitments and *regularity* accounts of the correctness of such commitments are species of one genus. They share a common strategy for naturalizing the different norms they address. In tracing the relation between them, it is helpful to keep in mind the basic case in which making an assertion consists in noninferentially applying a ground-level empirical concept in a particular situation. Regularity theories are attempts to determine the boundaries of concepts—which determine the difference between correct and incorrect application—by appealing to regularities or patterns in the actual applications of the concept and the dispositions to apply that concept that are exhibited by an individual or a community. The sort of *correctness* of application of concepts that such theories aim to explicate is what is assessed by judgments of the *truth* of the resulting assertion.<sup>8</sup> What it is for *A* to take *B*'s claim that

something is a porcupine to be a correct application of that concept is for *A* to take the claim to be true, that is, for *A* to endorse it, for *A* to undertake or acknowledge commitment to that same content. This is a different deontic attitude from *A*'s attributing to *B* *entitlement* to the commitment undertaken by that claim, and so to that application of the concept; different sorts of normative status are involved.

The concept of *reliability* in making a claim or applying a concept presupposes, rather than analyzes, such a notion of *correct* claiming or concept-application. For a reliable performer is just one who generally produces a correct performance; assessments of reliability are assessments of the probability of correctness. Thus the issue of reliability cannot be raised until the question of correctness that regularity theories address has been answered. Nonetheless, these different theories share an approach. For reliability theorists offer an account of entitlement that appeals only to patterns or regularities of *correct* claiming or application of concepts. The reliabilist idea is that *entitlement* to a particular claim or application of a concept—a derivative sort of correctness of claiming—can be understood entirely in terms of *dispositions* to produce *correct* performances of that kind. The regularist idea is that such correctness of claims or applications of concepts can be understood entirely in terms of *dispositions* (in some variants, those that are in some sense communal dispositions) to produce performances of that kind.

One of the major difficulties raised for the strategy of construing the correctness of discursive performances in terms of regularities or dispositions specified in nonnormative terms was the gerrymandering problem. There is no single pattern or regularity exhibited by any set of actual or virtual performances; where there is one, there are many—indeed an infinite number. No matter what a candidate performance whose correctness is at issue is like, and no matter what the history to which it must answer is like, there is *some* way of specifying the pattern exhibited by those prior performances so as to include the candidate as just what is required to continue that pattern “in the same way.” The attempt to identify the normative distinction between correct and incorrect performances with the naturalistic distinction between regular and irregular performances fails because no performance is simply irregular (even relative to a specified set of performances with respect to which its coregularity is to be assessed), and so none would be counted as incorrect by such a criterion. Appeal to regularity and irregularity can do normative explanatory work only if there is some way of *privileging* some regularities over others—some way, in other words, of saying what the *correct* regularity is. The problem of sorting performances into correct and incorrect is transformed by the regularist strategy into the problem of sorting regularities into the relevant and the irrelevant, the ones that *ought* to be taken account of in assessments of correctness, and those that ought not. From this vantage point, regularity theories appear as merely putting off the normative issue, moving the bump in the carpet around rather than smoothing it out.



## 2. Barn Facades

Reliability theories share with regularity theories the same fundamental strategy for explaining in naturalistic terms the normative statuses involved in discursive practice (although the phenomena they address are at different levels). So it might be expected that the possibility of gerrymandering would raise similar difficulties for a pure reliability strategy for construing entitlements to claims that it does for a pure regularity strategy for construing the correctness of those claims. This is indeed the case. A striking illustration of how the gerrymandering considerations get a grip on assessments of entitlement in terms of reliability is provided by Goldman's barn-facade example.<sup>9</sup> The example is forwarded to show the inadequacy of an account that seeks to ground the cognitive authority of noninferential reports exclusively in features of the causal chain leading from the reported state of affairs to the perceptual reporting of it. The leading idea of such causal theories is that a true belief, paradigmatically one acquired perceptually, counts as knowledge just in case it is caused in the right way by what it is about. To see that such a particular causal connection is not sufficient to make a true belief qualify as knowledge (and so cannot by itself perform the explanatory job assigned to the entitlement condition by the tripartite analysis), Goldman suggests comparing two cases that are alike as far as the causal chain leading to a claim is concerned, but unlike in the status of that claim as knowledge.

In each case the subject is in ideal circumstances for visual perception and is confronted by what is in fact a barn. In each the subject responds to the visible presence of the barn by confidently reporting the presence of the barn. The causal chains in each case are entirely standard, the barns reflecting light, which travels undisturbed to the subject's retina, and so on. Yet one of the subjects is, and the other is not, without knowing it located in Barn-Facade County. The local hobby in that county is building incredibly lifelike trompe l'oeil facades of barns. In this county, 99 percent of what appear to be barns are actually such facades. Each subject would in fact, if confronted with such a facade (and not alerted to the special practices of the natives), confidently report the presence of a barn. Goldman's plausible claim is that the claim of the subject who is *not* in Barn-Facade County *does* express knowledge (is a claim that subject is in the relevant sense entitled to), while the claim of the subject who *is* in Barn-Facade County does *not* express knowledge.

For the first point: the fact that there are *some* circumstances under which a subject could be fooled does not in general preclude the subject from having knowledge in the case in which that subject is not fooled. As Austin argued, the fact that it is possible to make a replica of a sparrow so cunningly contrived that I cannot tell it from the real thing does not mean that I cannot see a sparrow and know that it is a sparrow.<sup>10</sup> The mere possibility of hyperbolic doubt does not entitle anyone to it and does not undermine

entitlements in ordinary cases. For the second point: even though the subject who has the cognitive bad luck to be in Barn-Facade County is in fact looking at a barn, it is in an important sense just an *accident* that that is the case. The county is rife with perceptual situations in which the subject would with equal confidence and warrant *falsely* report the presence of a barn. Under these circumstances, the belief just *happens* to be true, and the subject should not be taken to *know* that there is a barn present.

Goldman claims first that this sort of example shows that one must look beyond the particular causal antecedents of a belief in order to determine its status as one the believer is entitled to in the sense relevant to assessments of knowledge. For in this case what distinguishes the two subjects is not the causal chains connecting them to the barns but only the incidence of barn facades in the vicinity, which is causally irrelevant to their perceptual transactions with the barns they are in fact looking at. His second claim is that the way in which that difference of causally irrelevant circumstance makes a difference to the assessment of entitlement and hence of knowledge can be understood in terms of the variable *reliability*, in those different circumstances, of the belief-forming mechanism that leads to the true belief in each case. The same differential responsive dispositions, the same noninferential reporting capacity, is in play in both cases.

The difference is that in Barn-Facade County it is not a reliable mechanism, while in the rest of the (largely barn-facadeless) world it is. How reliable a belief-forming mechanism is, how likely it is to yield a true claim, a correct application of a concept, depends on the circumstances in which it is exercised. My inability to tell sparrows from cunning duplicates does not disqualify me from being a reliable reporter of sparrows, so long as my environment is quite unlikely to confront me with a ringer. If such duplicates were to become common, the reliability of my differential responsive dispositions would degrade (and with it my capacity to acquire knowledge thereby in the cases where all goes well), even though the way in which that mechanism would respond to each possible case remained the same throughout. The probability of being correct in a particular case depends on the actual incidence of indistinguishable phonies. Thus the notion of reliability of belief-forming mechanisms provides just what is wanted to explain the barn-facade cases.

### 3. *Gerrymandering and the Problem of Reference Classes*

Goldman's argument is decisive against exclusively causal theories of knowledge, and it shows how assessments of reliability can function in assessments of entitlement—particularly entitlement to commitments acquired as a result of noninferential reporting capacities. But (though he does not make the point) it also underscores the possibility of gerrymandering, and hence the inadequacy of construing cognitive entitlement exclu-

sively in terms of reliability. In the case of regularity theories of the correctness of the application of a concept, it is the boundaries of the concepts that can be gerrymandered in such a way as to preclude assessments of irregularity, and hence of error. In the case of reliability theories of entitlement, it is rather the boundaries of the reference class with respect to which reliability is assessed that can be gerrymandered in such a way as to preclude assessments of unreliability, and hence of lack of entitlement.

Goldman's idea is that reliability is an objective affair, determined by the objective probability of a correct judgment, given one's circumstances. But such probabilities vary with the specification of those circumstances. Given a reference class of relevantly similar cases, frequencies of success define objective probabilities. The question remains how a privileged reference class is to be determined. What is the *correct* reference class with respect to which to assess such probabilities?

If the reference class is restricted to the actual case of the perceptual judgment that a barn is present, even in Barn-Facade County (since in the case being considered by hypothesis one is actually looking at a barn) the frequency of correct judgments is 1. So relative to that quite restricted reference class one is totally reliable. If the reference class is widened to the whole county, the frequency of correct judgments is reduced to 1 percent. So, relative to that less restricted reference class, one is quite unreliable. But since the customs of Barn-Facade County are quite parochial, the relative frequency of barn facades in the country as a whole is quite low. Relative to the nation as a whole, one is quite a reliable noninferential reporter of the presence of barns. Relative to the state, one's reliability will fall somewhere in between. One of the nice things about this example is that here the metaphor of *boundaries* is made concrete, and the difficulty of selecting the proper boundary is literally geographic.

Focusing on the relativity of reliability to decisions about where to draw these boundaries makes it evident that the question "Reliable or not?" is underdetermined in exactly the same way that the question "Regular or not?" is underdetermined. There are always *some* regularities that are being instantiated, and (in the case where the claim one is making is true) there are always *some* reference classes with respect to which one is reliable. Using these naturalistic notions to stand in for genuinely normative assessments works only relative to some way of privileging regularities or reference classes. The notions of regularity and reliability cannot do all the work by themselves. For the concept of regularity cannot discriminate between regularities, and that of reliability or probability of success relative to a reference class cannot discriminate between reference classes. In the sense in which, given a regularity, there is an objective matter of fact as to whether a further performance continues it, there is no objective matter of fact as to which of the various regularities exhibited by a given history of actual or dispositional performances is the right one to assess correctness with respect to. In the

sense in which, given a reference class of relevantly similar cases, there is an objective matter of fact as to what the probability that a certain skill exercised in those circumstances will yield a correct performance, there is no objective matter of fact as to which of the various possible reference classes to which the case in question might be assimilated is the right one to assess reliability with respect to. An objective or naturalistic theory of cognitive entitlement cannot be derived solely from considerations of reliability, any more than an objective or naturalistic theory of the correct application of concepts can be derived solely from considerations of regularity.

#### 4. *Taking or Treating as Reliable*

The general strategy of this work is to supply what is wanting in regularity theories of correct concept-application by appealing to the social (in an *I-thou* sense) and practical deontic attitudes of taking or treating a performance as correct or incorrect. The paradigm is taking or treating an assertion as correct in the sense of endorsing it, undertaking that commitment oneself, which is taking what it says to be true. It is these attitudes on the part of interpreters, of the deontic scorekeepers who attribute discursive commitments, that privilege some regularities over others and give a sense to the notion of *correct* use of expressions and so applications of concepts. The deontic *status* of *being* a correct application of a concept is to be understood in terms of the deontic *attitude* of *taking* or *treating* such an application *as* correct. That attitude (endorsing a claim, undertaking an assertional commitment) cannot be understood apart from its role in the essentially social practice of giving and asking for reasons, making and defending knowledge claims. The norms implicit in the application of concepts are social and perspectival, not (to begin with) objective and naturalistic.<sup>11</sup>

Regularity theories attempt to naturalize the normative status of correct claiming or concept-application. The countervailing idea pursued here is to explain that status by saying what it is for a performance to be taken or treated in practice as having such a significance. This is to focus on the deontic attitudes of acknowledging conceptual norms by attributing normative statuses and significances. Reliability theories attempt to naturalize the normative status of *entitlement* to the commitment undertaken by making a claim or applying a concept. The corresponding countervailing idea to be pursued here is accordingly to explain that status by saying what it is for a performance to be taken or treated in practice as having such a significance. This is to focus on the deontic attitudes that acknowledge that status and attribute that significance.

What in practice privileges some of the reference classes with respect to which reliability may be assessed over other such reference classes is the attitudes of those who attribute the commitment whose entitlement is in question. Each interpreter implicitly distinguishes between reference classes that are relevant and those that are irrelevant to the assessment of reliability,

and hence of entitlement to claims, by the circumstances under which that interpreter accords cognitive authority to those claims. The sort of authority in question here is not that acknowledged by the interpreter's own endorsement of the claim—that is, taking it to be correct in the sense of taking it to be *true* (which is the sense of correctness addressed by regularity theories rather than by reliability theories). The sort of authority in question is rather that of having an inheritable entitlement: the sort that supports successful deferrals by others (potentially including the interpreter). It is the scorekeeping social practices that actually govern the use of an expression (in particular the acknowledgment of entitlement to the commitments undertaken by its assertional use) that supply what is missing from pure reliability theories.

It is tempting, from the point of view of such theories, to think of the choice of reference class as a merely pragmatic matter—in a sense of 'pragmatic' that restricts it to what concerns the interests and goals of those performing speech acts. So it might be thought that for some purposes and in some contexts I should be counted as knowing that a sparrow is in front of me, even though I would believe that also if a sufficiently lifelike replica were there instead, while for other purposes and in other contexts (for instance where the stricter standards appropriate to discussions of principled skepticism are in force) I should not. No doubt there is such a variation in standards of entitlement depending on what is taken to turn on the issue, and it may be particularly acute in connection with the word 'know'. But the contribution made by interpreters (those who attribute commitments and entitlements to commitment) to the determination of the boundaries with respect to which reliability is assessed are not "merely pragmatic" in this sense.

They make a fundamental contribution to the semantic content of empirical concepts. Indeed, this is one of the situations in which traditional ways of distinguishing semantic from pragmatic concerns can be seen to be inapposite. In particular, as will appear, what an interpreter takes to be the circumstances under which an expression can appropriately be used in noninferential reports—that is, when interlocutors are *entitled* to commitments because the acknowledgment of the commitment arises through the exercise of dispositions to respond differentially to various aspects of their environment—is an important feature of the empirical content the interpreter associates with that expression. The sort of authority accorded to noninferential reports, and the way the model of assertional practice can be extended to incorporate it, is discussed further below.

### III. OBSERVATION REPORTS AND NONINFERENTIAL AUTHORITY

#### 1. *Knowledge, Entitlement, and Understanding*

The topic of reliability theories of cognitive entitlement was introduced in connection with the thought that once the notion of entitlement

or positive justificatory status that matters for attributions of knowledge has been broadened by the recognition that a belief, claim, or commitment can in this sense *be* justified without having *been* justified—that justifying is not the only way that status can be acquired—the way seems open to dispensing entirely with inferential justifying in explaining the deontic status of entitlement. It was then pointed out that identifying the entitlement of a commitment with its being the output of a reliable process has the same sorts of difficulties with gerrymandering that plague its relatives that identify the correctness of a claim or application of a concept with its being in accord with a regularity exhibited by other claimings or applications of the concept. But these difficulties concern only one way of following out the original thought about the in-principle dispensability of inferential justification in explaining the status of knowledge claims. The deep mistake involved in completely decoupling justifying as giving reasons from cognitive entitlement has to do rather with the sort of *understanding* that is presupposed by claims to and attributions of knowledge.

An assertion, even if true, is not taken to express knowledge unless the one making it *understands* the claim being made. A practical grasp of the significance of making the claim is inseparable from an appreciation of its role as possible reason for other claims, and as something that reasons can in turn be offered for. It is being caught up in this way in the game of giving and asking for reasons that makes a performance the undertaking of a commitment (the making of a claim) in the first place. Unless one accords one's own performance such a significance (treats it as a move in that game), one is not making a claim, not undertaking a commitment that is eligible for the status of knowledge.

It is on this basis that Sellars objects to construals of cognitive entitlement exclusively in terms of reliability. Reliable differential responsive dispositions are only a necessary condition for observational knowledge. Parrots and thermometers can have such dispositions and so can be used by us in effect as measuring instruments to acquire knowledge. But what *they* have is not knowledge. For they do not understand the significance of their responses; they do not take those responses as reasons for further claims; and they do not understand claims as potentially in need of reasons. To decouple entitlement from reason-giving entirely is to jettison the inferential articulation in virtue of which the performances and commitments one is entitled to can be understood as *propositionally* contentful. It is to discard precisely what makes responses, however reliably produced, have the significance of undertaking *discursive* commitments. What is left is not a *cognitive* affair at all.

The most serious objection to a pure reliability theory accordingly is presented not by the general Wittgensteinian strand of thought concerning the significance of gerrymandering for attempts to construe norms as regularities, which Kripke expounds so forcefully. The most serious objection to reliabilism stems rather from the more particular Sellarsian insight concern-

ing the essential role played by the specifically inferential articulation even of noninferential reports. Sellars insists first that for a performance elicited by a reliable differential responsive disposition to be a candidate for expressing knowledge, it must count as an *endorsement* by the reporter of some claim, as the undertaking of a *commitment*. Furthermore, he recognizes that the identity and content of such commitments depend on their role in inference and justification, in giving and asking for reasons. He sees further that being capable of endorsing a claim requires grasping the role of that claim in inference and justification—that the official tripartite analysis of knowledge implicitly involves *understanding*, as part of what is required for *belief*.

Unfortunately, motivated by these insights, Sellars stakes out far too strong an antirealist position concerning the role of inferential justifying in entitlement to claims to observational knowledge. There is accordingly a danger that where the various strands of thought are not carefully sorted out, distaste for the epistemological internalism about cognitive authority that Sellars endorses will obscure the important lessons that ought to be drawn from his account. His basic point is that a noninferential reporter must be “in the space of giving and asking for reasons,” in addition to having the right differential responsive dispositions. That space is, for Sellars as here, articulated by relations of *authority*, and to be in that space one must be able to recognize or acknowledge the authority of claims. It is, in other words, in virtue of one’s capacity to adopt practical deontic attitudes, to take or treat something as having cognitive authority, that one counts as moving in the space of giving and asking for reasons.

## 2. Sellars on the Authority of Noninferential Reports

In the passages (from “Empiricism and the Philosophy of Mind”) that follow, Sellars is concerned with the nature of the authority (or as he sometimes puts it, “credibility”) of noninferential reports (Carnap’s *Konstatierungen*), which are claims to observationally acquired knowledge. The distinctive feature of such reports is that “the credibility of such tokens as ‘express observations’ [is] a credibility which flows from tokens to types.”<sup>12</sup> This contrasts with the credibility of nonobservational claims such as “Dogs are mammals,” which is attached to tokens in virtue of their being of types that are credible. The reliability approach then recommends itself as offering a simple and natural account of the source and nature of the credibility of sentence tokenings that report empirical observations: “An overt or covert tokening of ‘This is green’ in the presence of a green item is a *Konstatierung* and expresses observational knowledge if and only if it is a manifestation of a tendency to produce overt or covert tokens of ‘This is green’—given a certain set—if and only if a green object is being looked at in standard conditions.”<sup>13</sup>

What does such reliability have to do with *authority*? “The first hurdle to

be jumped concerns the *authority* which, as I have emphasized, a sentence token must have in order that it may be said to express knowledge. Clearly, on this account the only thing that can remotely be supposed to constitute such authority is the fact that one can infer the presence of a green object from the fact that someone makes this report."<sup>14</sup> This is an important move. The authority of reliability consists in its underwriting a propriety of *inference* (what might be called "the reliability inference"). The noninferential undertaking of a commitment by a reliable reporter can *inferentially authorize* another to undertake a commitment with that content. To take or treat someone as a reliable reporter (in certain circumstances) is for a scorekeeper to endorse the propriety of the move from *attributing* to the reporter a noninferentially acquired doxastic commitment to the scorekeeper's *undertaking* of a corresponding commitment (and taking others to be similarly entitled). This notion of how reliability fits into the giving of reasons is the key to understanding the special sort of authority characteristic of noninferential reports, which in turn is essential to the notion of *empirically* contentful claims.

As is by now familiar, Sellars has already taken issue with the sort of foundationalism that sees empirical knowledge as an inferential superstructure raised on an autonomous noninferential base. The target of his criticism

is the idea that there is, indeed *must be*, a structure of particular matter of fact such that

(a) each fact can not only be noninferentially known to be the case but presupposes no other knowledge either of particular matter of fact, or of general truths; and

(b) the noninferential knowledge of facts belonging to this structure constitutes the ultimate court of appeals for all factual claims—particular and general—about the world.<sup>15</sup>

Sellars, inferentialist and antifoundationalist though he is, does not deny either (b) or the first half of (a). His quarrel is with the second half of (a).

There are particular instances of believing or being committed that are noninferential in the sense that their acquisition was not the conclusion of an inferential process. There are no beliefs or discursive commitments that are noninferential in that what is expressed by a sentence can be *understood* without mastering inferential relations that content stands in to others. So a bit of knowledge (belief) can, and indeed all of it does, presuppose other knowledge (belief), even though it is not inferred from that other knowledge or belief. This possibility was not seriously examined by the classical epistemological tradition. It is a certain hierarchical picture of *understanding* (at this level a necessary condition of believing) that Sellars rejects. He does *not* object to a hierarchical picture of empirical *justification*, once that has been suitably disentangled from bad foundationalism concerning understanding. His claim that the authority that accrues to noninferential reports in virtue



of their being the results of reliable reporting or belief-acquiring mechanisms is a broadly inferential authority is in no way inconsistent with understanding observational knowledge to be authoritative in virtue of the reliable *noninferential* differential responsive dispositions that produce it. Inference need not be involved in the process that leads to a tokening of 'This is green'; but it is involved in grasping the type of authority that such noninferentially produced tokenings have, and so in understanding such tokenings, and so in their being potential expressions of knowledge.

The question is just what the relation is between mastery of this inference and possession by a tokening of the sort of authority characteristic of ground-level observational knowledge. Sellars's claim is that "to be the expression of knowledge, a report must not only *have* authority, this authority must *in some sense* be recognized by the person whose report it is."<sup>16</sup> The notion of claims *having* cognitive authority is indeed intelligible only in connection with practical attitudes of *taking* or *treating* claims *as* having such authority. Sellars has suggested that the authority distinctive of observational knowledge should be understood in terms of the correctness of an inference, from the making of a report such as 'This is green' by one whose differential responsive dispositions are taken to be reliable to the undertaking of a commitment to the effect that there is something green there.

He concludes: "In other words, for a *Konstatierung* 'This is green' to express observational knowledge, not only must it be a *symptom* or *sign* of the presence of a green object in standard conditions, but the perceiver must know that tokens of 'This is green' *are* symptoms of the presence of green objects in conditions which are standard for visual perception."<sup>17</sup> So Sellars's view is that the reliable reporter can count as being entitled to a noninferentially acquired commitment, and so the assertion by which that commitment is acknowledged can be cognitively authoritative in licensing or entitling others by the standard assertional mechanism of communicative entitlement inheritance, *only* if the reporter can inferentially *justify* the noninferential claim. Such a justification consists precisely in exhibiting the inference whose premises are the reliability of differential responsive dispositions to make such claims and responsive elicitation of the claiming in question and whose conclusion is another tokening of the claim itself.

### 3. *Attributing Reliability Is Endorsing an Inference: An Inferentialist Middle Way between Justificatory Internalism and Reliabilist Externalism*

There are two problems with this conclusion. First, Sellars takes it that for the claim of the reliable observer to *be justified*, the observer must be able to *justify* it inferentially—to offer reasons by displaying premises from which it follows. Second, he assumes that such justification must involve explicit invocation of reliability, that is, that a claim of reliability

must be one of the premises. Thus Sellars claims that “observational knowledge of any particular fact, e.g. that this is green, presupposes that one knows general facts of the form *X is a reliable symptom of Y.*”<sup>18</sup>

This latter is an odd move for Sellars to make. He, after all, is the one who urged that material proprieties of inference not be everywhere traded in for suppressed premises. Why should it not be that reliability underwrites a propriety of inference, without the claim of reliability having to appear as a premise in the inferences so underwritten? Here Sellars overreaches himself. He is right to insist that reliability matters because it warrants inferences of a certain form. He is wrong to insist that this warranting must be understood in terms of endorsement of an explicit claim that can serve as a premise in inference.

The first claim also seems too strong. Securing entitlement to a claim need not always be assimilated to inferential justifying of the claim. The possibility of vindication of a commitment by deference rather than inference—by the invocation of communicational mechanisms depending on intracontent interpersonal testimony rather than intrapersonal intercontent justifying—shows that much. It would be a mistake to assimilate deferential entitlement inheritance to inferential entitlement inheritance by insisting that the one invoking the authority of another’s assertion be able to produce an explicit *argument* in which a claim as to the informant’s reliability would appear as a premise. Rather, deferring involves an *implicit* claim as to the reliability (more particularly the entitlement in this case) of the informant. One who accepts the deferring as successful (and so attributes the claimed entitlement to the one deferring) thereby implicitly endorses the propriety of a permissive inference from the informant’s claiming that *p* to *p*—which is just the inference that Sellars picks out as corresponding to reliability. But at the ground level, all of this can be made sense of as implicit in what is done in practice. Because it can, it is possible to explain the expressive role of the locutions that can be introduced at a later stage to make these attitudes explicit. Assimilating entitlement to the commitments acknowledged by noninferential reports, and therefore their authority, to that secured by explicit inferential justification, as Sellars does, is a mistake of the same sort.

A symmetrical mistake would be to assimilate the authority of noninferential reports to that of testimony, by understanding the invocation of such authority as a kind of deference to a “world-asserter.”<sup>19</sup> The structure of authority exhibited by noninferential reports is *sui generis*, to be reduced neither to that of inferential justification nor to that of testimony. These three are mutually irreducible—none can take over the function of any of the others. One of the primary explanatory aims of this work is to explain how commitments that are implicit in the fundamental practices that confer assertible conceptual content can eventually themselves be made explicit and assertible, expressed in a form in which reasons can be given and asked for them. The implicit attitudes that can in this way be explicitly expressed

once suitable vocabulary has been introduced include those involved in invocations and recognitions of the authority of both testimony and reports of observations. But the locutions that play these explicating roles can themselves be made intelligible only by understanding first the implicit structures they bring out into the assertional light of day.

In fact, Sellars's insight concerning the irreducible role played by inferential justification does not require insisting that noninferential reporters can be authoritatively entitled to their claims only if they can justify them. As just indicated, the inference from the noninferential undertaking of a commitment as the result of a reliable differential responsive disposition to endorsement of the claim thereby made may be implicit in practical attitudes, rather than explicit in claims offered as justifications for that endorsement. Nor is it necessary that the one who makes an observation report endorse the propriety of that inference, even in this implicit practical sense. Reliability may entitle the reporter to the knowledge claim, may qualify it as knowledge, even if the reporter does not even implicitly endorse the inference that is the practical acknowledgment of the authority of reliability. This is the primary insight that stands behind the justificatory *externalism* of reliability epistemologies, in contrast to the justificatory *internalism* Sellars exemplifies.

Suppose that Monique has been trained reliably to discriminate hornbeams by their leaves. As a result of the training, she is often disposed to respond to the visibility of leaves of the right sort by noninferentially reporting the presence of a hornbeam. She understands what it means to claim that something is a hornbeam and, in circumstances appropriate for such reports, actually comes to believe that there is a hornbeam present. She may still be uncertain of her discriminatory capacity long after she has in fact become reliable.<sup>20</sup> In such a situation she may have a true belief that there is hornbeam in front of her, yet be completely unable to justify that claim (for instance, by citing features distinctive of hornbeam leaves), and even deny that she is a reliable noninferential reporter of hornbeams.

Yet, the reliabilists point out, it can be entirely in order for one who does take her to be a reliable reporter of them, not only to come to believe that there is a hornbeam present on the basis of her report, but to cite her report (at least deferentially) as what warrants that belief. This is treating the claim as authoritative in just the way that is required for knowledge. Someone who thus takes her to be reliable can accordingly attribute to Monique the knowledge that there is a hornbeam in front of her, in spite of her protestations to the contrary. What makes her claim knowledge (according to the attributor) is the fact of her reliability (according to the attributor), regardless of her attitudes toward that reliability.<sup>21</sup> The status of her claim as knowledge is accordingly external to her attitudes—not only because of the truth condition on knowledge, but also because of the entitlement condition. Sellars is committed to withholding the attribution of knowledge in the absence of the

candidate knower's capacity to justify the claim, and so is committed to disagreeing with reliabilists about examples like this. Yet on this point the reliabilists are surely correct.

Sellars, however, is right that for a reliably elicited differential response to be a candidate for knowledge, the one making the knowledge claim must be in the space of reasons, must be capable of understanding the claim, and so must have some grip on its role in reasoning, hence on its use as a premise and conclusion of inferential justifications. Requiring this general capacity, of course, falls short of requiring that on each occasion the reporter must be able to justify the claim for it to count as the expression of observational knowledge. Furthermore, while reliabilism about cognitive entitlement and so cognitive authority is clearly correct that knowledge can be attributed even where the one to whom it is attributed cannot demonstrate entitlement to the claim inferentially, by providing a justification that appeals to other claims the putative knower endorses, it does not follow from this observation that reliability by itself is enough for entitlement and cognitive authority, apart from all consideration of attitudes of taking or treating the knower as reliable, as a thoroughgoing externalism about entitlement would have it. Sellars is also right to insist that attributions of knowledge require not just reliability but at least implicit endorsement of the inference that is the practical acknowledgment of reliability—the inference namely from the occurrence of a report, or the noninferential undertaking of the commitment such a report expresses, to the endorsement of the claim.

Where Sellars is wrong, as the sort of example just considered shows, is in thinking that the one who endorses this inference must be the one who *undertakes* the claim to observational knowledge. It has been pointed out that attributing knowledge is a hybrid deontic attitude involving not only the attribution of commitments but the undertaking of them. Not only does the attributor of knowledge take the candidate knower to endorse a claim; the attributor also must endorse that claim.

It is likewise the *attributor* of observational knowledge who must attribute reliability to the knower. Attributing such reliability is endorsing exactly the general form of permissive inference that Sellars points to—treating as appropriate the inference from the noninferential undertaking of a commitment (of the right sort, and in circumstances of the right sort) by the observer to the endorsement by others of the claim so elicited. Taking a report that is the outcome of a particular differential responsive disposition as entitling others to the claim (for instance accepting as entitling their deferrals to the reporter on such issues) just *is* treating the reporter in practice as reliable about such matters. Monique need not, *pace* Sellars, take herself to be a reliable reporter of hornbeams in order for her to count as knowing observationally that there is a hornbeam in front of her. But the one who attributes such knowledge must take her to be reliable. And adopting that practical attitude is endorsing the pattern of permissive inference that connects the

*attribution* by others to Monique of a noninferentially acquired belief about the visible presence of hornbeams with their *undertaking* of a commitment to the visible presence of hornbeams in Monique's vicinity.

Just as the truth condition on knowledge requires that the attributor of knowledge undertake, as well as attribute, commitment to the content of the knowledge claim, so satisfying the entitlement condition by mere reliability requires that the attributor of knowledge undertake (but not necessarily attribute) commitment to the propriety of the reliability inference. Where the language is rich enough to include the expressive resources necessary to make the reliability inference explicit (conditionals and 'claims that . . .' or 'believes that . . .'), attributors of knowledge can be challenged and called on to defend their endorsement of the conditional "If Monique claims (sincerely, responsively, and in appropriate conditions) that a hornbeam is visibly present, then (probably) a hornbeam is present." At this point, reliability could be invoked to *justify* the belief that there is a hornbeam present. But this is a sophisticated, late-coming possibility, built on the implicit acknowledgments already described. So full-blown reliabilist externalism about cognitive entitlement is mistaken in ignoring the necessity of such inferential attitudes on the part of attributors of knowledge, while full-blown Sellarsian internalism about cognitive entitlement is mistaken in insisting that the knower must have such attitudes. These are complementary ways of misunderstanding the essentially social structure of the cognitive deontic attitudes, in terms of which the status of a claim as knowledge must be understood.

#### 4. *Observational Knowledge and Empirical Conceptual Content*

The noninferential authority possessed by claims issuing from the exercise of reliable differential responsive dispositions—although reducible neither to the sort of interpersonal authority invoked by deferring nor to the sort of intercontent authority invoked by inferring—is not fundamental in the way that those structures of authority are. In the model of assertional practice that has been put on the table, communication and justification are two aspects of the game of giving and asking for reasons; neither is intelligible except in the context of the other. They are intelligible, however, in the absence of noninferential responsive authority. Practices that do not involve according any knowledge claims the significance of observation reports can nonetheless be understood as instituting specifically *assertional* significances, and so as conferring specifically *propositional* contents.

What is missing from such practices is claims with *empirical* content. Discourse recognizable as mathematical can be like this: reasons are given and demanded; claims communicated, challenged, and justified; and regresses of entitlement inheritance halted by appeal to axioms, free moves that anyone is treated as entitled to at any point in the conversation. Our

discourse is not in general like this, however, and the sorts of contents our claims have cannot be conferred by assertional practices that do not acknowledge some claims as having empirical authority stemming from their status as reports of observations. Indeed, it is essential to the contents of the ordinary concepts in terms of which we conduct our lives that they stand in inferential relations both to the acknowledgments of commitments resulting from what Sellars calls "language entry transitions," in perception, and to the acknowledgments of commitments that result in what Sellars calls "language exit transitions," in intentional action.<sup>22</sup> (The contribution of the latter *practical* empirical structure is discussed in the second half of this chapter. Attention is restricted here to the *cognitive* empirical structure.)

The practical significance characteristic of claims to observational knowledge is best understood in terms of the role they play in the default-and-challenge structure of entitlement. Noninferential reports can function as unjustified justifiers: claimings that are treated as having a defeasible default status as entitled. Properly made claims to observational or perceptual knowledge can accordingly provide entitlements that can then be inherited either inferentially or communicatively. So observation provides regress-stoppers, and in this sense a foundation for empirical knowledge. This is what stands behind Sellars's endorsement of the claim (quoted above) that "noninferential knowledge of facts . . . constitutes the ultimate court of appeals for all factual claims—particular and general—about the world."<sup>23</sup>

Default entitlements are of two sorts, depending on whether the entitlement attaches to a commitment in virtue of the type it instantiates or in virtue of the circumstances in which it is tokened.<sup>24</sup> There are sentence types that would require a great deal of work for one to get into a position to challenge, such as "Red is a color," "There have been black dogs," "Lightning frequently precedes thunder," and similar commonplaces. These are treated as "free moves" by the members of our speech community—they are available to just about anyone any time to use as premises, to assert unchallenged. Noninferential reports, by contrast, have their default entitlement status as a result of the way in which the report tokening, or the particular acknowledging of the commitment that would be expressed by such a tokening, is elicited through the exercise of a reliable differential responsive reporting disposition. Treating such a claim as one the reporter is entitled to involves an implicit commitment on the part of the attributor to the actual circumstances being among those in which the reporter is responsively reliable concerning the sort of matters reported.

There will typically be some sorts of reports such that under appropriate reporting conditions (the same for all), essentially all the members of the linguistic community are reliable. Almost anyone can, under suitable circumstances, tell whether it is a warm day out or whether the marble one is holding is approximately round. Other sorts of reports involve not only more specialized circumstances but specialized training. Particle physicists are

trained reliably to respond noninferentially to the presence of mu-mesons in a bubble chamber by reporting the presence of mu-mesons. Not all of us can do this reliably. It is only someone who is taken not only to be looking at a bubble chamber but also to be properly trained to be reliable about these matters (and who has the right sort of collateral beliefs) whose reports will be accorded noninferential entitlement and the corresponding authority. As Quine says, what is observable varies from community to community.<sup>25</sup> He understands the status of being an observation report for a community (perhaps a proper subset of the whole linguistic community) in terms of what that community can agree on under concurrent stimulation, that is, in the same standard reporting circumstances.

The authority of noninferential reports requires the collaboration of both dimensions into which Sellars analyzes them: not only that they arise from the exercise of reliable differential responsive dispositions but that the response is to endorse a claim, to acknowledge a commitment, with a certain content. What makes it a mu-meson that the physicist is reporting rather than the hooked vapor trail that also forms part of the reliably covarying chain of events culminating in the report is to be understood not in terms of the differential responsive dispositions but in terms of the inferential role of the claim being made. The consequences that can be inferred from the presence of a mu-meson are quite different from those that can be inferred from the presence of a hooked vapor trail, covariant and concomitant though these phenomena may be. For instance, mu-mesons are much smaller, and move much faster, than the vapor trails they produce (see further at 7.1.6 below). As Quine argues further, it is important to understand that under the appropriate circumstances, which include the presence of a bubble chamber or similar device, and for the right community of observers, mu-mesons are literally observable—noninferentially reportable in much the same sense in which red things are for the rest of us. It is a mistake to think that what is really noninferentially observed is only the vapor trail and that the presence of mu-mesons is only *inferred*. Such an inference can be made, and learning to make it might be part of the training process that leads to becoming a reliable observer of mu-mesons (in bubble chambers). But coming to be disposed reliably to respond to the vapor trail, and *hence* to the presence of mu-mesons, by asserting or acknowledging a commitment to the presence of a mu-meson is learning to *observe* mu-mesons, to report them *noninferentially*. And this is so even if one is not totally reliable, in that there are circumstances in which one would *mistakenly* report the presence of a mu-meson *because* of the presence in the chamber of a vapor trail indistinguishable from those one has learned to respond to noninferentially by reporting mu-mesons (just as the fact that one can be fooled by a cunning replica does not preclude one from seeing a sparrow in the cases where one is not being fooled).

The claim is, then, that one is *directly* observing mu-mesons, in the sense

of *noninferentially* coming to be aware of them, to make claims about them, to know something about them—rather than indirectly, inferentially coming to a conclusion about mu-mesons on the basis of an *inference* (perhaps unconscious or implicit) from the presence of a vapor trail with a certain shape. This claim may seem implausible in light of the common practice of retreating, under certain suitable sorts of challenge, from the claim that a mu-meson is present to the claim that a hooked vapor trail is present. Is not such a retreat to be understood as relinquishing commitment to an inference, and therefore to its conclusion, while continuing to defend its (genuinely noninferential) premise? No.

Such cases ought to be understood as retreats (given a credible challenge to the effect that this might be one of the cases in which the exercise of a generally reliable capacity nonetheless leads one astray) to a claim that is *safer*. Being safer in this sense, however, is not a matter of withdrawing endorsement of an inference. One retreats to a different report with respect to which one is *more* reliable, as measured for instance by percentage of correct differential responses in the relevant circumstances, or by the same percentage of correct responses within a wider range of circumstances, or by the size of the community that does not share esoteric theoretical beliefs but does share the differential responsive disposition and corresponding capacity to make noninferential reports. Doing this can amount to offering an inferential justification of the original noninferential belief, by explaining how one was able to see it. An analogous case would be explaining that there was a mirror, not apparent to the audience assessing the authority of the claim, but apparent to the reporter, in order to explain how one was noninferentially able to report something that a challenger has pointed out is around a corner and so ought to be invisible.

The possibility of such an inferential justification of a claim on the basis of a safer claim does not show that the original claim should be understood as itself the product of a process of inference, any more than the capacity of sophisticated reporters to offer justifications of their claims to observational knowledge by citing their reliability and appealing to the reliability inference shows that their original claim was arrived at as a result of an inference from that premise. Nor does the fact that the capacity to make certain sorts of noninferential reports depends on collateral beliefs show that those reports are really inferences from something more basic, together with those collateral beliefs. One must have many beliefs about mu-mesons in order to be able to understand and so to make *any* claims about them, noninferential or otherwise. That does not preclude one from coming to be able to observe them.

The basis of observational knowledge, then, is that it should be possible to train individuals reliably to respond differentially to features of their environment by acknowledging doxastic commitments. Those commit-



ments are inferentially related to others that not only play inferential roles but also are themselves appropriately elicited noninferentially by features of the environment. These cross-connections put constraints on endorsements of inferences relating expressions whose circumstances of appropriate application include noninferential ones. Both oranges and orange things can be noninferentially reported, so someone who reports the presence of an orange after tasting but not seeing it and then infers from its being an orange to its being orange in color is liable to be challenged by another who is in a position to report it noninferentially as purple. For the commitment entitlement to which was acquired noninferentially is incompatible with that entitlement to which was acquired inferentially. Either the identification of the orange by taste, which formed the premise of the inference, or the identification of its color as purple might itself bear further challenge and investigation; but if these stand up, the reliability of the inference from being an orange to being orange in color will be impugned. In this way the possession of noninferential circumstances of appropriate application of some concepts imbues them with empirical content—recognizable as conceptual content in virtue of its inferential articulation and as empirical in virtue of its dependence on the noninferential acquisition of commitments to those contents (and of entitlements to those commitments).

Similarly, the inferences from circumstances to consequences of application (which are implicit in conceptual contents) are subject to empirical criticism in virtue of inferential connections among the contents of commitments that can be acquired noninferentially. So it may happen that one uses the term 'acid' in such a way that a substance's tasting sour is a sufficient condition for applying it, and that it will turn litmus paper red is a necessary consequence of applying it. Finding a substance that both tastes sour and turns litmus paper blue shows that such a concept is inadequate. Conceptual contents can accordingly be criticized, groomed, and developed empirically in a way parallel to the sort of Socratic process discussed in Chapter 2. In virtue of their inferential connections to concepts that can be used to make reports, even purely theoretical concepts (those whose *only* circumstances of appropriate application are inferential) inherit empirical content and have the inferences they are involved with constrained by the commitments and entitlements actually thrown up by what is responded to noninferentially. That the reliable differential responsive dispositions underlying this structure of noninferential authority are dispositions to acquire *commitments* and *entitlements* to those commitments, that is, to alter deontic status, means that the practices they appear in must include corresponding practical deontic *attitudes*. Something practitioners can do must be the taking or treating of performances as having the significance of noninferential reports, the recognition of the status of some claims as deriving their entitlements from their being expressions of reliable differential responsive dispositions to ac-

knowledge commitments. For one cannot make sense of normative significance, even the normative significance of reliability, apart from consideration of its uptake or attribution.

### 5. *Attributing Observational Entitlement*

It is straightforward to extend the model of assertional practice as outlined so far to encompass the structure of authority in virtue of which claims can have and be treated as having the significance of noninferential reports. What is primarily required is to say what it is for one interlocutor to attribute noninferential or observational authority to the claim of another, thereby recognizing or acknowledging it as having a special sort of entitlement. The authority involved is entitlement heritable by the usual intrapersonal intercontent inferential and interpersonal intracontent communicational pathways. What is distinctive of observational authority is that such authority is accorded to particular *tokenings* of acknowledged commitments (rather than to their types) and the way in which that authority depends on a special combination of content-based and person-based features. For the imputed reliability of an observer varies from content to content within each observer, and from observer to observer—someone who is taken to be able reliably and noninferentially to discriminate mu-mesons in bubble chambers may not be taken to be able to do so for '52 Pontiacs in traffic.

Furthermore, if the topic is fixed (the concepts used in the reports being assessed are specified), imputed reliability still varies depending on circumstances. For each particular observable, there will be an associated set of appropriate circumstances of reporting, according to the one attributing reliability and so observational authority. The authority of a reliable reporter is conditional on the obtaining of those appropriate circumstances. Even a lookout who is in general a reliable reporter of whales must be facing in the direction of what is being reported, cannot see well in the direction of a horizon-hugging sun, is less reliable if there are large walruses about, and so on. These appropriate circumstances of reporting, associated with the observable content (and perhaps the individual reporter), figure as commitments *undertaken* by the one attributing or assessing the responsive authority of a claiming.

So associated with each sort of noninferential authority a given interlocutor grants to another (the product of a person and a kind of content, for example, reports of the presence of whales, or mu-mesons), there is a set of enabling conditions (looking in the right direction, looking in a bubble chamber) and a set of defeating conditions (presence of many walruses in the vicinity, physicist drunk and woozy), and it is the interaction of these, according to the commitments undertaken by the one assessing noninferential authority, that determine whether responsive entitlement is attributed or

not. If the assessor undertakes commitment to a suitable range of the enabling reporting conditions and does not undertake commitment to any of the defeating reporting conditions, the reporter's claim is treated as having a default status as entitled. In this way the empirical authority of some attributed commitments, on the basis of implicit inferential acknowledgment of reliability under suitable circumstances, is distinguished from the type-based default status of "Red is a color" and "There have been black dogs," which do not exhibit a similar relativity to person, content, and the enviroing conditions as they are taken to be by the assessor. Observational authority is accordingly another hybrid deontic status: attributing it involves not only *attributing* commitments and entitlements but also *undertaking* or *acknowledging* them by endorsing reliability inferences.

Once the attitude of taking or treating someone's performance as having the significance of a noninferential report whose authority is grounded in the local and conditional reliability of the observer is in place, it is possible to introduce a type of performance that is the claiming of or petitioning for such authority by an observer. A certain sort of noise or gesture (perhaps a shrug) can come to have the significance of *invoking* observational authority. Then if a report is challenged, it can be vindicated (its entitlement demonstrated) by invoking observational authority rather than by deferral or inferential justification. But there is no strict need for practices encompassing empirical conceptual contents to include a speech-act kind with this significance. It is enough if interlocutors sometimes accord such authority, and thus take the commitments acknowledged by noninferential reports in some circumstances to be vindicated (implicitly, according to the assessor's attitudes) by the fact of the reporter's reliability. Where there is such a speech act, it would implicitly mean something like "I *see* it (for example, that it is red)." An explicit assertion to this effect can be introduced as well, but just how will not be clear until Chapter 8, where pragmatically explicitating locutions such as 'believe that' and 'claim that' are officially introduced into the model of assertional practices. The significance of an invocation of observational authority does not depend on any assertionally explicit content that the invocation might have. (Compare Wittgenstein's suggestion that if challenged to say how one knows that the thing in front of one is red, one might say simply, "I speak English.")

### 6. Expressions of Belief That Are Not Claims to Knowledge

The account here of assertions as claims to *knowledge* turns on the implicit obligation to vindicate the commitment undertaken by demonstrating one's *entitlement* to it. The foregoing discussion of observation and reliability focused on the importance for the *attributor* of observational knowledge claims of implicitly attributing reliability. Adopting that attitude requires endorsing the inference from the attribution of a noninferentially

responsively elicited acknowledgment of a *commitment* (under suitable circumstances and for a qualified observer) to the attribution of *entitlement* to that commitment. It is in the context of concern with entitlement to assertional commitments that the complaint was levied against reliabilists that they ignore the inferentially articulated attitude in which recognition or attribution of entitlement consists. The corresponding objection to Sellars was that, while appreciating the significance of that hybrid practical deontic attitude, he inappropriately insists that the reliability inference it involves be endorsed by the one *making* the observation rather than the one attributing or assessing it.

But this concern with entitlement can seem out of place in a discussion of a sort of discursive commitment that is intended to do the sort of explanatory work characteristically performed by a notion of *belief*. If belief is to be understood in the first instance as the state or status expressed by assertional speech acts, it seems wrong to treat assertions as also involving a claim to *knowledge*. For expressing a belief and claiming to know are different.

When an idiom is developed to the point that it has the expressive resources provided by the English locutions 'believes that' and 'knows that'—which make the pragmatic status being attributed or undertaken explicit as part of the *content* of what is claimed—it becomes possible to say of someone else, "He believes that Arnauld did not write *The Art of Thinking*, but he does not know it." The case has already been considered where what is expressed is the attitude of a scorekeeper who attributes commitment to a claim but does not endorse the attributed claim—that is, does not take it to be true. It is also possible, however, to distinguish expressions of mere belief from claims to knowledge in the first-person case, in which the claim *is* being endorsed or taken-true. In such cases, the social-perspectival distinction between attributions of knowledge and attributions of belief cannot get a grip.

For although undertaking an assertional commitment *is* taking-true the claim, a difference can arise precisely over the issue of entitlement or justification. The attribution of knowledge may be withheld by a scorekeeper who attributes a commitment without attributing a corresponding entitlement. Indeed, sometimes we make claims while fully aware that they may legitimately be challenged and that we are not in a position to vindicate them by demonstrating our entitlement to them. This is the implicit attitude that becomes assertionally explicit in claims such as "I believe that Arnauld did not write *The Art of Thinking*, but I don't claim to know it." For this sort of reservation can concern not the truth of the belief but my capacity to justify it. I may continue to take the claim to be true, to endorse it, to acknowledge the commitment it expresses, and yet not be prepared to shoulder the justificatory burden associated with a knowledge claim. This might be because I have forgotten the source of my conviction, or it might be an expression of a claim's having a ground-level status for me as an unjustified justifier that I do not take to be widely shared—I just believe that people with beards cannot be trusted, or that house cats are dangerous.

The speech acts that express such attitudes are what might be called *bare* assertions,<sup>26</sup> corresponding to *mere* beliefs, without the implicit claim to entitlement that is demonstrable should someone become entitled to challenge it (paradigmatically by expressing an entitled commitment to a claim incompatible with it). Does not the possibility of such bare expressions of commitment without claim of entitlement, of conviction without warrant, show that it is a mistake to understand claims on the model of claims to *knowledge*? No. Such claims are intelligible only as exceptions against a background of practices in which claims typically have the significance of claims whose authority *is* redeemable by demonstration of warrant. The possibility of bare assertion is parasitic on the possibility of assertions that implicitly involve undertaking a conditional task-responsibility to demonstrate the asserter's entitlement to the commitments undertaken by the performance of speech acts of that kind.

For bare assertions and the commitments they express would be completely idle if they could not figure as premises in inference and could not be passed along in communication. This is to say that bare assertions involve something of the *authority* of full-blooded assertions, while disavowing the corresponding *responsibility*. Yet that authority (licensing inferences by the asserter to commitments with other contents and the undertaking of commitments with the same contents by other interlocutors) makes sense only in a context in which inferential and deferential invocation of such authority can be demanded. What assertions are *for* is justifying other assertions. To accept someone's bare assertion is to take it to be a claim from which conclusions can be drawn. But giving reasons presupposes the possibility of asking for them, or at least the possibility that claims often stand in need of reasons. A game of giving and asking for reasons cannot consist exclusively in the exchange of speech acts that are accorded the significance of bare assertions. Within the broader context of full-blooded assertions (which do involve the demonstration of entitlement by inference and deference), however, it is possible to make sense of treating some claims as having the significance of bare assertions. Assertional commitments essentially involve the dimension of entitlement. Assertions are paradigmatically knowledge claims, and the sort of belief they express is unintelligible except in relation to the possibility of assessing beliefs for their status as knowledge, as warranted and true.

#### IV. RATIONAL AGENCY

##### 1. *Methodological Constraints on the Conception of Practical Rationality*

Beliefs make a difference both to what we say and to what we do. They manifest themselves both linguistically, in assertions, and practically,

in actions. A basic criterion of adequacy for any theoretical account of this fundamental sort of intentional state is that it explain both of these ways in which beliefs can be expressed in behavior, and the relation between them. The methodologically parsimonious idea that one or the other of them ought to be accorded explanatory priority is the motivation common to both of what Stalnaker (in the rough-and-ready botanization alluded to in the previous chapter) distinguishes as the “linguistic” and the “pragmatic” approaches to intentionality. Theories of the sort he calls “linguistic” construe believing by analogy to claiming: as a kind of inner asserting of sentences. They are accordingly obliged, first, to explain assertions without appeal to their role as expressions of belief and, second, to explain the norms that determine the role of belief in rational agency in terms of the proprieties that govern the public use of sentences. Theories of the sort he calls “pragmatic” (such as the one Stalnaker himself endorses), in contrast, take the role of belief in intentional action to be primary. They then owe both a nonlinguistic explanation of rational agency and an account of speech acts, paradigmatically assertion, in terms of intentional states so understood.

The approach pursued here takes belief to be intelligible only in the context of social-linguistic practice. But it is a relational, rather than a reductive, linguistic theory. Although doxastic commitment (the sort of deontic status corresponding to the intentional state of belief) cannot be made sense of apart from the possibility of expressing such commitments by performing speech acts that have the significance of assertions, neither can assertional significance be made sense of without reference to the commitments such speech acts undertake and acknowledge. As regards asserting and believing, the theory is even-handed; it accords explanatory priority to neither one. It nonetheless deserves to be called a *linguistic* account of intentionality (in a sense broader than Stalnaker's<sup>27</sup>) inasmuch as it does accord explanatory priority to the linguistic manifestation of belief in assertion over its practical manifestation in action.

As Dennett and Davidson have emphasized, attributing propositionally contentful intentional states such as beliefs to a creature is taking it to be *rational*. Thus Kant uses the rubrics of *theoretical* and *practical* rationality to distinguish the sort of normative competence manifested in giving and asking for reasons for *claims* or *judgments* from the sort of normative competence manifested in giving and asking for reasons for *actions*—judgments and actions being picked out precisely as the sorts of things reasons can be given for and for which reasons can be asked. Being rational is understood here generically as being able to play the game of giving and asking for reasons, which is to engage in a specifically *linguistic* social practice. For one cannot give reasons unless one can make claims. Doing so requires mastery of the normative dimension of inference: a practical grasp of the notion of *right* reasoning, of the distinction between correct and incorrect inference. Assessing performances as correct or incorrect is adopting normative atti-

tudes that are intelligible only in a context of interpersonal scorekeeping—even though in such a context it is possible to make sense both of self-assessment and of assessments of objective correctness, for which no one's scorekeeping attitudes are counted as decisively authoritative.<sup>28</sup> First-person deliberation is the internalization of such third-person assessment.

To take this line—identifying the rationality that qualifies us as sapient with being a player in the normative game of offering and assessing, producing and consuming reasons—is to deny two widely held reductive conceptions of rationality: one that identifies rationality with *logical* competence, and another that identifies it with prudence or *instrumental* competence. All parties can agree that to be rational is to distinguish good inferences from bad inferences. The disagreement concerns whether 'good inference' in this formula can be restricted to *logically* good inferences, or again to *instrumentally* good inferences—ones whose correctness is determined by their utility in satisfying desires or maximizing preferences. Logical competence is mastery of the use of locutions by means of which inferential proprieties are made *explicit* as the contents of claims. This theoretical ability to codify practices as principles accordingly presupposes prior practical mastery both of *implicit* inferential proprieties and of the use of the ordinary, nonlogical claims they articulate and govern. When this expressive role of logical vocabulary is appreciated, the identification of rationality in general with its manifestation as logical manipulation is unmasked as another form of the intellectualism that insists on discerning a propositionally explicit principle underlying every implicit propriety of practice—a form of platonism whose remedy is a complementary pragmatism.

Identifying rationality in general with the sort of instrumental rationality manifested in rational agency also inverts the proper order of explanation.<sup>29</sup> For the propositional contents of the intentional states appealed to in practical reasoning presuppose assertional-inferential proprieties, and hence linguistic social practices. (Though to say this is not to deny that proprieties of practical inference also contribute to the propositional contents of the states and expressions caught up in them.) To make out this claim it is necessary to say something about the practical reasoning that is implicitly attributed in interpretations of individuals as rational agents. In particular, just as it was shown how the capacity for logical reasoning is to be made intelligible in terms of (as the explicitation of) a conceptually prior capacity for nonlogical reasoning, it must also be shown how the capacity for practical reasoning incorporates and depends upon a conceptually prior capacity to give reasons for claims, rather than for actions.

Only an account of assertion of the sort introduced in Chapter 3 leaves room for the pursuit of such an order of explanation. Everyone ought to agree that asserting is putting forward a sentence as true. Following Davidson's lead, it has been suggested that distinguishing practical attitudes as taking or treating something *as true* requires a specifically linguistic social context of

mutual interpretation—that is, attribution of doxastic commitments, of the sort that has been elaborated as assertional-inferential scorekeeping. The next chapter develops the idea that this principle is best exploited by starting with an antecedent notion of assertional significance and then moving via that principle to an understanding of what is involved in talk of truth.

Commitment to understanding rational agency in terms of linguistic practice, rather than the other way around, strongly constrains the construal of the putting-forward portion of the principle that asserting is putting forward a sentence as true. For the claim that our first grip on the paradigmatic intentional state of belief (taking-true) is as what is expressed by assertions—rather than as what makes certain nonlinguistic performances intelligible in a way that is made explicit by exhibiting a piece of practical reasoning—is evidently incompatible with understanding asserting instrumentally, as a means intentionally adopted by a rational agent in order to achieve certain desired ends. If the linguistic practice of making and assessing claims (the game of giving claims as reasons and demanding reasons for claims) is an essential element of the context required to make sense of the notion of *propositional* intentional content (assertible, believable contents, which in English can be made explicit by the use of ‘that’ clauses), then what has been called “agent semantics” is not entitled to the conceptual raw materials it employs. In particular, one may not appeal to the *intentions* of the asserter—for instance intentions to say something true, or to make the audience believe that what is said is true, or to make the audience believe that it is uttered with the intention of saying something true or of engendering the corresponding beliefs. For what one is attributing can be identified as intentions to bring about various states of affairs only in virtue of the role such states play in a larger practical whole—one that includes the possibility of attributions of beliefs that the corresponding states of affairs obtain.

Less obviously, this order of explanation also precludes appeal to *conventions*, at least as commonly understood. The influential account offered by Lewis, for instance, takes a convention to be a social regularity that is sustained in a special way by the beliefs, intentions, and desires of the parties to the convention.<sup>30</sup> They are required not only to conform to the regularity but, among other conditions, to believe that others do so, to conform themselves because of that belief, to prefer that everyone conform, and to believe that everyone else has such beliefs and preferences. The present view is that on such a construal of convention, as Davidson concludes, “philosophers who make convention a necessary element in language have the matter backwards. The truth is rather that language is a condition for having conventions.”<sup>31</sup> Construing the putting-forward bit of the principle that for a sentence to have assertional significance is for it to be put forward as true in terms of social conventions rather than individual intentions is also not an acceptable move according to this explanatory strategy. Certainly conventions of the sort that Dummett tries out—conventions to the effect that one



is to be understood as trying to utter sentences only with the intention of uttering true ones—are of no avail in the context of these explanatory commitments.<sup>32</sup> That is why it was necessary to move beyond explaining asserting in terms of *intentions* or *conventions* to explaining it instead in terms of *practices*, which themselves can be explained without appeal to intentions or conventions.

The next task is to show how that account of practices can be extended so as to encompass deontic statuses corresponding to the other sorts of intentional states that figure in the giving of reasons for *nonlinguistic* performances: the intentions and desires that play an essential role in the *practical* reasoning implicitly attributed by interpretations of individuals as rational *agents*. The aim is to provide a broadly Kantian account of the will as a *rational* faculty. By exploiting the analogy between discursive *entry* transitions in *perception* and discursive *exit* transitions in *action*, the rational will can be understood as no more philosophically mysterious than our capacity to notice barns or red things. A scorekeeping account can pick out performances (largely nonlinguistic ones) as intentional (under some specification) and hence as actions (under any specification) insofar as they are expressions of deontic attitudes—acknowledgments of a certain kind of commitment. Practical reasoning can then be understood as leading to performances with this sort of deontic significance. And on that basis, the expressive role of distinctively *normative* vocabulary can be specified in terms of its role in making explicit the endorsement of patterns of practical reasoning.

## 2. Acting and Perceiving

The general claim is that there are two species of discursive commitment: the cognitive and the practical. Acknowledging commitments of these two sorts is adopting deontic attitudes that correspond to the intentional states of *believing* and *intending*, respectively. A practical commitment is a commitment to *act*. The content of such a practical commitment is to *making-true* a claim. These commitments and their contents are intelligible only in a context that includes also the *taking-true* of claims. For it is in terms of such assertional taking-true that the success of actions, the fulfillment of practical commitments, must be understood. The category of cognitive discursive commitments accordingly enjoys a certain explanatory priority over that of practical discursive commitment. Each is essentially something that reasons can be given for and for which reasons can be asked, and one cannot give reasons unless one can acknowledge doxastic commitments by making claims.

The practical dimension of discursive practice can be understood by exploiting two ideas. The first is that practical commitments are like doxastic commitments in being essentially inferentially articulated. They stand in

inferential relations both among themselves and to doxastic commitments. The second idea is that the noninferential relations between acknowledgments of practical commitments and states of affairs brought about by intentional *action* can be understood by analogy to the noninferential relations between acknowledgments of doxastic commitments and the states of affairs that bring them about through conceptually contentful *perception*. The causal dimension of acting for reasons—acknowledging practical commitments by acting on them—involves the exercise of reliable differential responsive skills on the *output* side of the game of giving and asking for reasons, just as perception does on the *input* side. Elaborating the first idea involves examining the sense in which practical reasons are *reasons*; elaborating the second idea involves examining the sense in which practical reasons are *causes*.

Adding practical commitments to the model of discursive practice enriches the propositional contents that such practice can be understood to confer on states and their expressions in a way analogous to the enrichment provided by including the empirical authority of observationally acquired doxastic commitments. In each case the general category of assertional commitments and their contents can be understood in advance of the enrichment. The three structures of authority that the model of assertion, as presented thus far, comprises are mutually irreducible, but not all are equally fundamental. The inferential authority invoked by justification and the testimonial authority invoked by deference are intelligible apart from the default authority of noninferential reports; but inferential and deferential practice are two sides of one coin, apart from which the authority of noninferential reports is not intelligible. Thus *empirical* content represents an enrichment of the generic sort of propositional content specifiable in abstraction from the contribution of observation. Similarly, *practical* content represents an enrichment of the generic sort of propositional content specifiable in abstraction from the contribution of action. The empirical and practical involvements of claims—even those that are purely theoretical in the sense that they are only inferentially connected to claims that have direct empirical and practical significance—make a fundamental contribution to their contents. Only a model that incorporates both of these not purely inferential dimensions of discursive articulation has any prospect of generating propositional contents that resemble those expressed by the declarative sentences of natural languages.

The best way to understand the place of action in the deontic model of discursive practice is to exploit the analogy between action and perception. Sellars divides the “moves” that can be made in a language game into three kinds: intralinguistic moves, language entry moves, and language exit moves.<sup>33</sup> The first kind consists of inferential moves. These are moves in which a position within the language game (paradigmatically the endorsement of a claim) is responded to by the adoption of another such position.

The second kind consists of noninferential reports of observations. These are moves in which a nonlinguistic situation is responded to by the adoption of a position within the language game (paradigmatically the endorsement of a claim). The third kind consists of deliberate actions. These are moves in which a position within the language game (for instance, endorsement of a plan) is responded to by bringing about a nonlinguistic situation.

Following Sellars's lead, language entry moves have been analyzed in the first three sections of this chapter in terms of two components in their content: their inferential articulation and their noninferential elicitation. In virtue of the former they are *conceptually* contentful, and in virtue of the latter they are *empirically* contentful. These components and their interaction have been elaborated here in the idiom of deontic scorekeeping, into which Sellars's framework has been transposed. In that idiom, noninferential reports count as *entries* in the sense that they are responses that consist in changes of deontic scorekeeping attitude, elicited by stimuli that do not themselves consist in changes of deontic score. As such they contrast with inferential moves, in which an alteration of deontic attitude—for instance the undertaking or attributing of a commitment—has as a scorekeeping consequence another alteration of deontic attitude. The language or discursive scorekeeping *exits* (intentional *actions*) are to be understood by analogy to these *entries* (perceptual *observations*). In action, alterations of deontic attitude, specifically acknowledgments of practical commitments, serve as stimuli eliciting nonlinguistic performances.

Observation depends on reliable dispositions to respond differentially to states of affairs of various kinds by acknowledging certain sorts of commitments—that is, by adopting deontic attitudes and so changing the score. A competent observer under suitable circumstances responds to the visible presence of a red ball by coming to acknowledge a commitment to the claim that there is a red ball present. The content of the commitment responsively undertaken is jointly determined by the chain of reliably covarying events that culminates in its acquisition and by its inferential connection to other contents (including those empirical conceptual contents that themselves incorporate a responsive observational component). Action depends on reliable dispositions to respond differentially to the acknowledging of certain sorts of commitments (the adoption of deontic attitudes and consequent change of score) by bringing about various kinds of states of affairs.<sup>34</sup> A competent agent under suitable circumstances responds to the acquisition of a commitment to flip the light switch by flipping the light switch. The content of the commitment so expressed is jointly determined by the chain of reliably covarying events that its acknowledgment initiates and by its inferential connection to other contents (including both other contents that themselves incorporate a practical component and those empirical conceptual contents that incorporate a responsive observational component).

In any given situation, interlocutors can be taught to be reliable noninfer-

ential reporters of only certain sorts of circumstances. Unaided by special instruments, we cannot reliably discriminate the presence of X rays, and we cannot tell automobiles that will at some point in the future be painted green from those that will not; we can reliably discriminate the presence of loud noises, and we can tell automobiles that are now painted green from those that are not. Similarly, interlocutors can be taught to be reliable performers of only certain kinds of acts. Unaided by special tools, we cannot reliably produce X rays, and we cannot make an automobile have been painted green some time in the past; we can reliably produce loud noises and paint automobiles green. What can be noninferentially reported varies from reporter to reporter and from situation to situation. Only a properly trained physicist can noninferentially observe the presence of a mu-meson, and then only with a bubble chamber; only a properly trained pianist can noninferentially produce a performance of the Moonlight Sonata, and then only with a piano.

Observation requires reliable responsive dispositions to acquire acknowledged commitments, while action requires reliable responsive dispositions to fulfill acknowledged commitments. Reliability in the first case concerns the relation between the state of affairs responded to and the content of the commitment acknowledged. Reliability in the second case concerns the relation between the content of the commitment acknowledged and the state of affairs brought about. In each case, assessments of reliability require some independent access to the eliciting or the elicited state of affairs—assessments of the *truth* of the claim the perceiver has noninferentially come to make and of the *success* of the performance the agent has noninferentially come to produce. Attributions of reliability consist in endorsements of score-keeping inferences *from* commitments *attributed* to reporters or agents *to* commitments *undertaken* by the attributor of reliability (commitments concerning the state of affairs reported or produced). Thus my noninferentially acquired doxastic commitment to the effect that there is a red thing in front of me is, under appropriate conditions, a good reason for others inferentially to acquire a doxastic commitment to the effect that there is a red thing in front of me. My acknowledging a practical commitment to the effect that I will raise my arm in the next minute is, under appropriate conditions, a good reason for others to undertake a doxastic commitment to the effect that I will raise my arm in the next minute.

In observation, the elicited commitment-acknowledgment is an attitude toward a *doxastic* discursive deontic status. In action, the eliciting commitment-acknowledgment is an attitude toward a *practical* discursive deontic status. The first sort of attitude corresponds to believing or taking-true—in *one* sense of believing, namely the causally relevant sense that depends on what one would *acknowledge* commitment to, not the ideal sense in which if  $p$  entails  $q$ , then believing that  $p$  is believing that  $q$ , whether one knows it or not. The second sort of attitude corresponds to intending or making-true—in *one* sense of intending, namely the causally relevant sense that depends on what one would *acknowledge* commitment to, not the ideal sense in

which if doing *A* entails doing *B*, then intending to do *A* is intending to do *B*, whether one knows it or not.

The wider ideal senses of 'believe' and 'intend' correspond to the deontic *statuses* of doxastic and practical commitment, rather than to the deontic *attitudes* of acknowledging them. These senses are to be understood in terms of the fundamental scorekeeping principle that *undertaking* a commitment (to begin with, by acknowledging it) licenses others to *attribute* it, and the attributions that are thereby authorized can outrun what one is disposed to acknowledge. The only function of the concept *deontic status* in the idiom in which the model of discursive practice is formulated is its use in keeping score. It is a creature of the activity of scorekeeping on deontic attitudes; deontic statuses figure only as the objects of attitudes, as what is undertaken and attributed.

Understanding practical discursive commitments (commitments to act) is accordingly a matter of understanding their pragmatic significance: the way they depend on and influence the deontic score interlocutors keep by acquiring and relinquishing attitudes toward their own and others' deontic statuses. Practical commitments, like doxastic or assertional commitments (including noninferentially acquired empirical ones), are discursive or conceptually contentful commitments in virtue of the *inferential* articulation of their pragmatic significance. The scorekeeping significance of practical commitments is analogous to that of doxastic commitments—indeed the inferential and incompatibility relations that the contents of practical commitments stand in are largely inherited from those of corresponding doxastic commitments, except for their role in the sort of practical reasoning that connects them inferentially with doxastic commitments proper. Thus one practical commitment can have others as consequences; a commitment to drive to the airport today entails a commitment to go to the airport today, because the inference from 'X drives to *s*' to 'X goes to *s*' preserves *doxastic* commitments. In the same way, one practical commitment can be incompatible with another, as are a commitment to drive to the airport today and a commitment to spend the day snoozing in a hammock under a shade tree—again because of the incompatibility of the corresponding doxastic commitment contents.

The instrumental inferences corresponding to the principle "Who wills the end wills the means," like inferences generally, come in two flavors: committive and permissive. Some instrumental inferences (those whose premises specify goals one is committed to and whose conclusions specify the necessary means to those ends) are also commitment-preserving. If cutting down a tree is the only way to get across the ravine, then undertaking or attributing a commitment to getting across the ravine has as a scorekeeping consequence undertaking or attributing a commitment to cutting down a tree. But some means-end reasoning is permissive in nature; there may be more than one way to skin a cat.

Inferences whose premises express commitments to secure certain ends and whose conclusions express sufficient (but not necessary) means to those

ends are entitlement-preserving rather than commitment-preserving. One who is entitled to a practical commitment to secure an end is entitled thereby to a practical commitment to performances that would (according to the one whose scorekeeping is being elaborated) bring about that end—in the absence (as is always the proviso with permissive inferences) of collateral commitments incompatible with such a commitment. Entitlement to a practical commitment to achieve some end may simultaneously entitle one to each of a set of mutually incompatible alternative means; entitlement to a commitment to cross the ravine may instrumentally entitle one to cut down the tree at the edge of the ravine, and it may entitle one to anchor a rope bridge to the top of that tree, even though doing one of these things precludes doing the other. In the same way, permissive inferential relations (paradigmatically inductive ones) among the contents of doxastic commitments can result in entitlement to each of a set of incompatible conclusions. In each case, choosing one, committing oneself to a conclusion or a means, relinquishes entitlements to those incompatible with it. For incompatibility is a relation involving both deontic statuses: two contents, whether doxastic or practical, are incompatible in case commitment to one precludes entitlement to the other. Entitlement to both without commitment to either is not ruled out. Neither, of course, is commitment to both; making this possibility straightforwardly intelligible is one of the cardinal advantages deontic normative construals of belief have over causal-functional ones.

### *3. Asymmetries between Practical and Doxastic Discursive Commitments*

It is in their relations to their corresponding entitlements that practical discursive commitments differ most markedly from doxastic discursive commitments. The significance of undertaking a doxastic commitment, paradigmatically through its overt acknowledgment by assertion, was explained in terms of the interactions between the coordinate dimensions of authority and responsibility. The responsibility involved is to vindicate the commitment, by demonstrating or displaying one's entitlement to it, if it is brought into question by a suitable challenge (an incompatible assertion with an equal, *prima facie* claim to entitlement). Default entitlements aside, this responsibility can be discharged by appeal to the authority of other doxastic commitments; credentials for the commitment are secured by displaying its entitlement as inherited from that attached to other commitments.

The authority of the commitments undertaken by assertion exhibits a dual structure, corresponding to two different sorts of routes by which entitlement can be passed on for use in discharging the responsibility associated with other commitments. On the one hand, a doxastic commitment to which one interlocutor is entitled licenses further commitments (with different

contents) by that same interlocutor. These are its *inferential* consequences (committive and permissive). This sort of authority is invoked to vindicate those consequential commitments by presenting a *justification*, which appeals to the authorizing claims as premises. On the other hand, a doxastic commitment to which an interlocutor is entitled licenses further commitments with the same content, by other interlocutors. This is its authority as *testimony*. It is invoked to vindicate the commitments it authorizes, by *deferral* to the one whose testimony is relied upon.

The first way in which the structure governing the attribution of entitlements to practical discursive commitments differs from that governing the attribution of entitlements to doxastic ones is that there is nothing corresponding to the authority of testimony in the practical case. The issue of entitlement can arise for practical commitments, as for all discursive commitments. But the (conditional) responsibility to vindicate such commitments is, in the practical case, exclusively a *justificatory* responsibility. Default entitlements aside, it is only by exhibiting a piece of reasoning having as its conclusion the practical commitment in question that entitlement to such commitments can in general be demonstrated or secured.

This feature of the deontic scorekeeping model of discursive commitments reflects a fundamental asymmetry between expressing a belief by making a claim and expressing it by performing an action. What I take-true I thereby, *ceteris paribus*, authorize you to take-true. Though there can be various complications about the transfer of title (because of differences in collateral beliefs), in general what serve me as good reasons for belief can serve you also as good reasons for that same belief. What I (seek to) make-true, however, I do not thereby in general authorize you also to (seek to) make-true. What serve me as good reasons for action may or may not be available to you as good reasons for action, even bracketing differences in collateral beliefs. For you and I may have quite different ends, subscribe to different values, occupy different social roles, be subject to different norms. That I have good reasons to drive to the airport today does not imply that you do. If you form a similar intention, you cannot in general show that you are entitled to it by deferring to me ("Well, *he's* going"). Only some kinds of reasons that entitle me to an intention and action are automatically available to you. You might be in a position to make the same argument I can, but if so, that in general is independent of my being in a position to use that line of thought; there is no general (even defeasible) presumption of heritability.<sup>35</sup>

Committing oneself to a claim is putting it forward as *true*, and this means as something that everyone in some sense *ought* to believe (even though some unfortunates will for various reasons not be in a position to do so and need not be blameworthy for that failure). Committing oneself to a course of action need not be like this. It need not (though in special cases it can) involve putting it forward as something that everyone else ought to do (even subject to the recognition that some unfortunates will for various reasons not

be in a position to appreciate this, and need not be blameworthy for that failure). Some kinds of reasons for actions, paradigmatically moral ones, have a permissive or committive force that is independent of interpersonal differences. But reasons for action in general do not have this kind of force. What an agent has reason to do can depend on what that agent wants (or on what institutional role that agent occupies). Differences among agents as to desires and preferences (or institutional roles) need not have the significance of indications of normative failures. Whenever two believers disagree, a diagnosis of error or ignorance is appropriate for at least one of them. Though agents with differing practical commitments can also be criticized on the grounds of error and ignorance, mere difference of desire or preference is not sufficient in general to make them liable to such criticism.

We come with different bodies, and that by itself ensures that we will have different desires; what is good for my digestion may not be good for yours; my reason to avoid peppers need be no reason for you to avoid peppers. Our different bodies give us different perceptual perspectives on the world as well, but belief as taking-true incorporates an implicit norm of commonality—that we should pool our resources, attempt to overcome the error and ignorance that distinguish our different sets of doxastic commitments, and aim at a common set of beliefs that are equally good for all. Talk about belief as involving an implicit commitment to the Truth as One, the same for all believers, is a colorful way of talking about the role of testimony and challenge in the authority structure of doxastic commitment—about the way in which entitlements can be inherited by others and undercut by the incompatible commitments they become entitled to. The Good is not in the same way One, at least not if the focus is widened from specifically moral reasons for action to reasons for action generally, so as to include prudential and institutional goods. Desires and preferences can supply reasons for actions (can entitle agents to practical commitments) in the sense of ‘entitle’ that corresponds to that at stake in the discussion of doxastic commitments, and desires and preferences can vary from individual to individual. That there is no implicit normative commitment that plays the same role with respect to desire (and therefore intention and action in general) that truth plays with respect to belief consists simply in the absence (in the structure according to which entitlements to practical commitments are inherited) of anything corresponding to the interpersonal dimension of testimony and vindication by deferral.

It is of course possible to add an interpersonal dimension of practical authority as a superstructure to the basic game of giving and asking for reasons for actions. Where within a certain sphere of practical activity the performance of one individual licenses or compels performances by others, there exists an authority relation of superior to subordinate. In a practice in which reasons can be given (and so asked for) at all—that is, a linguistic practice, one in which some performances are accorded the significance of



assertions—the authorizing performances can be speech acts with the significance of imperatives and permissives. The superior issues an order, which specifies what the subordinate is obliged to do (what the subordinate thereby acquires a commitment to do), by displaying the assertion that must be made-true, the assertible content of the doxastic commitment that anyone must be entitled to undertake (perhaps observationally) upon completion of the task. Or the superior offers permission in the form of a licence, which specifies what the subordinate is entitled to do by displaying the assertion that can be made-true.

This sort of practical authority structure is like that of testimony in some ways. When the issue of the agent's entitlement to a practical commitment (perhaps claimed implicitly by deliberate action) is raised, rather than the entitlement being inherited from reasons that could be cited for it—either by the agent in terms of doxastic and practical commitments undertaken or by other scorekeeping assessors in terms of doxastic and practical commitments attributed to the agent—that entitlement can be inherited from the superior who ordered or permitted it. Such authority can be invoked by deferring to the issuer of the command or license (a mode of vindication codified in the legal doctrine of *respondeat superior*). So besides intrapersonal entitlement inheritance invoked by inference, there can be a mode of interpersonal entitlement inheritance invoked by deference, in the practical as well as the doxastic case.

There are many disanalogies between these two cases as well, however. First, the licensing is restricted as to subject matter and the interlocutors involved, to those situations in which a prior superior/subordinate authority relation has been established. The employer can authorize or compel only certain sorts of performances, and only on the part of certain individuals. Perhaps this difference does not go very deep. The limitation is characteristic of a society in which such authority relations are established and limited by explicit contracts. In a society based on status rather than contract, the superior/subordinate relations are fixed once and for all in advance, appearing as part of the nature of things, and need not be restricted as to subject matter at all. Furthermore, such restrictions can arise, *de facto* or even *de jure*, in the case of assertional authority as well: the teacher of secret doctrines may not authorize their repetition or answer for them except to favored students. And there are what amount to hierarchies of assertional authority regarding technical topics such as mu-mesons and quarter horses.

One difference that does go deep, however, is an asymmetry between the authorizing performance and the authorized performances, in the case of commands and the issuing of licenses. The asserter licenses members of the audience to perform speech acts with just the same content and significance as the original assertion. They are authorized to authorize others in the same sense in which they are authorized. Assertion, at least as it is construed in the ideal Sprachspiel presented here, is an egalitarian practice in a sense in

which commanding and giving permission is not. Only in very special cases does the practical license one is given authorize the further issuance of such licenses; only in very special cases does the command one is given compel or permit one to offer such commands to others.

The structure of default entitlement and calling to account by challenging entitlements to action are similarly asymmetric in the case of interpersonal practical authority structured by superiors and subordinates. There is in general nothing corresponding to assessments of reliability underlying the default authority of superiors (though analogs exist for special cases). Entitlement to challenges must similarly be relativized to superior/subordinate relations, if commands and licenses are to have any significance. The point of rehearsing these asymmetries is just that the fundamental differences between doxastic and practical structures of authority and entitlement inheritance remain even in the case where the normatively significant social status of individuals as superior or subordinate is widely or universally instituted by the attitudes of those keeping score on commitments and entitlements. If the asymmetries characteristic of superior/subordinate relations are removed, making interpersonal practical authority look more like assertional authority as here conceived, nothing remotely resembling the issuing of orders or the giving of permission results.

In the interests of simplicity, the deontic scorekeeping model of assertional significance has been talked about as though assertional authority is always made universally available throughout the community and is always universally recognized. Where testimony has this sort of catholic significance, the community can be thought of as engaged in the search for a single common body of truths, for anyone's entitlement to any claim is open to challenge from any quarter. Doxastic practice need not be so monolithic, of course. There may be many subcommunities, distinguished precisely by what sorts of authority they acknowledge, and so what sorts of challenges to entitlements they take to be in order. Specialists may recognize the authority only of other specialists. Members of one speech community may be divided into competing schools of thought on various topics and may not recognize the entitlements or therefore the challenges of those from other groups, as regards claims concerning those topics. Yet within those subcommunities it is essential that the authority granted by an assertion include a reassertion license—a license to do just what the asserter did. This feature makes it possible for the claims of one interlocutor to have the significance of challenges to the claims of another.

The importance of this structure is particularly evident in the case of empirical practice, for it is by testimony that observations by one interlocutor can be assessed and adjusted by confrontation with the observations of another. The notion of entitlement to a doxastic commitment depends on the in-principle heritability of interpersonal authority. Because an assertion that would be defended by appeal to testimony can have just the same

entitlement status as one that would be defended by providing a justification or by the invocation of noninferential responsive reliability (observational prowess), the credentials of each claim do not need to be traced back to their source before it can be treated as having the significance of a *prima facie* challenge to incompatible claims. This is a basic feature of the assertional default-and-challenge structure. That a status or performance whose entitlement is inherited from another should have just the same authority as the status or performance that authorized it (according to the subcommunity that recognizes such authority) accordingly distinguishes doxastic discursive commitments from practical discursive commitments. For as has been pointed out, if subordinates have the same authority as their superiors in virtue of being commanded or licensed by them, the entitlement of a superior to issue a command would be subject to challenge by commands issued by subordinates, not just by other superiors, and the hypothesized asymmetry between superior and subordinate would disappear.

## V. PRACTICAL REASONING: INFERENCES FROM DOXASTIC TO PRACTICAL COMMITMENTS

### 1. *Acting for Reasons and Acting Intentionally*

Giving and asking for reasons for *actions* is possible only in the context of practices of giving and asking for reasons generally—that is, of practices of making and defending *claims*. The structure of entitlement characteristic of practical discursive commitments is not autonomous but presupposes that of doxastic ones. This dependence appears in two ways. On the side of the circumstances of acquisition of practical deontic statuses, it appears in the role of practical reasoning; practical reasoning requires the availability of doxastic commitments as premises. On the side of the consequences of acquisition of practical deontic statuses, it appears in the essential role that propositional (= assertible) contents play in specifying conditions of success—that is, what counts as fulfilling a commitment to act. With regard to this latter role, it has already been pointed out that practical commitments inherit some of their inferential relations from the propositional contents that specify their conditions of success. If doxastic commitment to *p* has as a scorekeeping consequence doxastic commitment to *q*, then a practical commitment to make-true *p* has as a scorekeeping consequence a practical commitment to make-true *q*. Understanding what one has committed oneself to by undertaking a practical commitment to bring it about that *p* accordingly requires mastery of the inferential role *p* plays in doxastic discursive practice.

The relation between doxastic and practical commitments that is most important for extending the deontic scorekeeping account to include both

species of discursive status, however, is that exhibited in practical reasoning. Because interpersonal inheritance of entitlements is not an essential part of the scorekeeping structure that institutes practical deontic discursive statuses, the conditional responsibility to demonstrate entitlement that is part of undertaking a commitment to act is a specifically *justificatory* responsibility. Justifying a practical commitment is exhibiting a suitable piece of practical reasoning in which it figures as the conclusion. It is in terms of practical inferences that we give reasons for action, make our own and each other's conduct intelligible, exhibit it as rational. Practical reasoning accordingly forms the core of intentional explanations of nonlinguistic deportment.

In what follows, an abstract account of practical reasoning is sketched in the deontic scorekeeping terms familiar from the treatment of theoretical reasoning concerning doxastic commitments.<sup>36</sup> The explanatory framework in which the notion of practical reasoning is to function is the Kantian one, according to which to treat a performance as an *action* is to treat it as something for which it is in principle appropriate to demand a *reason*. Not everything an agent does is an action. If I am walking along the top of a cliff and stumble and fall off, stumbling and accelerating at 32 feet per second per second are both things I do (in the sense that they are bits of my behavior), but they are not actions of mine; walking and grabbing a bush as I topple over the edge are. Actions are the things agents do *intentionally*. In the terms to be employed here, acting intentionally is noninferentially producing a performance that either is the acknowledgment of a practical commitment (in the case of intentions in action) or results from exercising a reliable differential disposition to respond to such an acknowledgment (in the case of prior intentions). The acknowledgment of the practical commitment can be thought of as the intention with which the performance is produced.

One can act *with* a reason, but unintentionally (for instance in a case in which one is unaware of the commitments that supply the reason that an attributor might cite). But only what is done intentionally can be done *for* a reason—though one can act intentionally but without a reason.<sup>37</sup> Only rational beings can be agents, but there are such things as irrational actions: for instance where one acts intentionally, but on impulse rather than according to what one has reason to do. In the deontic framework, such irrational actions are *intentional* in that they are acknowledgments of practical *commitments* (or arise from the exercise of reliable noninferential dispositions to respond differentially to them), and they are irrational in that the practical commitment in question is not one the agent is *entitled* to by a good practical inference from premises that agent is committed and entitled to—either because one has *no* reason or because one has an overriding reason to do something incompatible with what one in fact does. Since to be so entitled requires having a *reason* for performing the action, practical commitments, and therefore actions (intentional performances), are attributed only to those who are in the space of giving and asking for reasons—that is, to those who are (treated as) rational.

Undertaking any discursive commitment involves a conditional responsibility to demonstrate entitlement to it. In the case of practical commitments this takes the form of a specifically justificatory responsibility. Only against the background of a general capacity to comprehend and fulfill such a justificatory responsibility—to assess and produce reasons for practical commitments—can what one does have the significance of an acknowledgment of a practical commitment, that is, the significance of acquiring or expressing an intention. Given such a general capacity or status as rational, however, one can in particular cases undertake practical commitments to which one is not entitled by reasons, and so act irrationally. Intentional but irrational actions are perfectly intelligible within the deontic framework, in the same way and for the same reasons that, on the side of doxastic rather than practical discursive commitments, incompatible beliefs are—namely as commitments lacking the corresponding entitlements. These phenomena cause explanatory difficulties for other sorts of accounts (for instance those that construe intentional states exclusively in terms of causal-functional role), difficulties that simply do not arise when those states are construed in terms of deontic statuses instituted by scorekeeping attitudes.

To be entitled to a practical commitment is to have suitable reasons for it. Practical inferences—as distinct from the doxastic inferences that have been considered in previous chapters—are those that have practical commitments as their conclusions.<sup>38</sup> Reasons for such commitments, and hence for the actions elicited by the acknowledgment of such commitments, are the premises of good practical inferences. It has already been pointed out that intentions can serve as reasons for other intentions—the intention to bring it about that  $p$  serving as a reason for intending to bring it about that  $q$  if that  $q$  is true is necessary or sufficient for bringing it about that  $p$ . What about reasons for commitments to act that are not themselves commitments to act? Facts, as acknowledged in doxastic commitments, can provide reasons for practical commitments.

## 2. *Three Patterns of Practical Reasoning*

Consider the following three bits of practical reasoning:

- ( $\alpha$ ) Only opening my umbrella will keep me dry,  
so I shall open my umbrella.
- ( $\beta$ ) I am a bank employee going to work,  
so I shall wear a necktie.
- ( $\gamma$ ) Repeating the gossip would harm someone, to no purpose,  
so I shall not repeat the gossip.

'Shall' is used here to express the significance of the conclusion as the acknowledging of a practical commitment.<sup>39</sup> The corresponding doxastic commitment would be acknowledged by a standard assertion using 'will'.

The role of the speech acts performed by uttering sentences of this form can be understood in scorekeeping terms from their fundamental pragmatic significance as acknowledging a practical commitment, together with the inferentially articulated content that results from combining its involvement in the inferences deriving from the corresponding 'will' statements with its involvement in the sort of basic practical inferences of which ( $\alpha$ ), ( $\beta$ ), and ( $\gamma$ ) are examples.<sup>40</sup>

There are two ways to think about inferences like these, which move from doxastic premises to practical conclusions, from beliefs to intentions. What is perhaps the standard approach is that taken by Davidson.<sup>41</sup> He defines a *primary* reason as the pair of a belief and what he calls a pro-attitude. He allows that sometimes one or the other is cited by itself as a reason, but insists: "In order to understand how a reason of any kind rationalizes an action it is necessary and sufficient that we see, at least in essential outline, how to construct a primary reason."

In other words, inferences such as those exhibited by ( $\alpha$ ), ( $\beta$ ), and ( $\gamma$ ) are *enthymemes*, in which a premise necessary for the correctness of the inference has been suppressed or omitted. In the first inference, what is missing is some such premise as:

(a) Let me stay dry,

an expression of a desire, preference, or pro-attitude that would be explicitly self-ascribed by something like "I desire (prefer) to stay dry."<sup>42</sup> The second inference might be underwritten by something like:

(b) Bank employees are obliged (required) to wear neckties.

In the third case, the suppressed premise is something such as:

(c) It is wrong to (one ought not) harm anyone to no purpose.

Each of these supplies the missing pro-attitude required to make the premises into primary reasons.

As appears in these examples, the notion of pro-attitude encompasses not only wants, desires, and preferences but also more general evaluative attitudes. This assimilation represents an important insight, which will be exploited below. In fact, Davidson thinks that all pro-attitudes are expressed by sentences that are in a broad sense normative or evaluative.

There is no short proof that evaluative sentences express desires and other pro-attitudes in the way that the sentence "Snow is white" expresses the belief that snow is white. But the following considerations will perhaps help show what is involved. If someone who knows English says honestly "Snow is white," then he believes snow is white. If my thesis is correct, someone who says honestly "It is desirable that I stop smoking," has some pro-attitude towards his stopping smoking.

He feels some inclination to do it; in fact he will do it if nothing stands in his way, he knows how, and he has no contrary values or desires. Given this assumption, it is reasonable to generalize: if explicit value judgments represent pro-attitudes, all pro-attitudes may be expressed by value judgments that are at least implicit.<sup>43</sup>

Davidson thinks of evaluative expressions as expressing something like desires, but the connection can equally well be exploited in the other direction. For conversely, one who desires or prefers  $p$  to  $q$  (say desires that one eat pears rather than that one eat peaches) thereby attaches some value to  $p$  over  $q$ , takes  $p$  to be preferable to, or more desirable than,  $q$ . What is important is to see that normative claims and expressions of desire and preference are species of a genus defined by the role they play in completing primary reasons.

Pro-attitudes must be included in primary reasons, on this account, to bridge the gap between what one believes and what one decides to do. My preference to stay dry makes my belief that I can stay dry only by opening the umbrella relevant to the practical issue of whether to open the umbrella. The fact that bank employees are obliged to wear neckties makes my working at the bank relevant to the practical issue of whether to wear a necktie. And the negative value of causing pointless harm (the fact that it is wrong or that one ought not to do it) makes the consequence of gossiping relevant to the practical issue of whether to gossip. In the context of different pro-attitudes, those same beliefs would provide reasons for quite different intentions and actions.

### 3. Normative Vocabulary Makes Explicit Material Proprieties of Practical Reasoning

There is another way of construing the relation between  $(\alpha)$ ,  $(\beta)$ , and  $(\gamma)$ , on the one hand, and  $(a)$ ,  $(b)$ , and  $(c)$ , on the other hand. That relation could be modeled on the relation between *materially* good inferences and the *conditionals* whose addition as premises would turn them into *formally* (logically) good inferences. In that case (considered in Section IV of Chapter 2), it turned out to be a fruitful strategy to consider the apparently enthymematic inferences as in order just as they stood, and to treat the conditionals not as suppressed *premises* but as making *explicit* (expressing in the form of a claim) what is *implicit* in the endorsement of the inferences. Part of the payoff from considering things this way around is an understanding of the expressive role played by conditionals; they can be understood as making *inferential* commitments propositionally explicit (= assertible). What makes that approach possible is an account of proprieties of inference as deontic social statuses instituted by scorekeeping attitudes, so that commitment to a material propriety of inference can be understood in terms of what it is to

take or treat an inference as correct in keeping score by attributing and acknowledging assertional commitments.

There is no bar to understanding  $(\alpha)$ ,  $(\beta)$ , and  $(\gamma)$  as materially good inferences in this sense. The fact that endorsement of claims incompatible with (a), (b), and (c) would void these inferences does not show that they function as suppressed premises—any more than the fact that endorsement of  $p \rightarrow \sim q$  would void the inference from  $p$  to  $q$  shows that the conditional  $p \rightarrow q$  is a suppressed premise in the material inference from  $p$  to  $q$ . The claims (a), (b), and (c) might, like  $p \rightarrow q$ , be understood rather as codifying material-practical inferential commitments. The payoff from doing so would be making it possible to understand the expressive role played by the broadly evaluative words (such as 'prefer', 'obliged', and 'ought') used to express these pro-attitudes, in a way analogous to the understanding suggested for conditionals. According to such an account, although the sort of practical inference instanced by  $(\alpha)$ ,  $(\beta)$ , and  $(\gamma)$  does not need supplementation to be correct, in a language with sufficient expressive resources it is possible to make the inferential commitments that are implicit in endorsing such inferences explicit in the form of claims.

The benefits of doing so are the familiar benefits of propositional explicitness: once expressed in the form of claims, these commitments are themselves subject to challenge and justification, rather than simply being accepted or rejected. Two interlocutors who disagree about the correctness of an inference such as  $(\gamma)$  can now argue about whether (c) is true, challenge entitlement to such a claim, and offer counterclaims to it. A new venue is opened up for resolving disagreements about what follows from what, about which claims rationalize which actions. Davidson's view (transposed into the deontic idiom) that in order to see how a doxastic commitment can rationalize a practical commitment we must be able to see ("at least in essential outline") how to construct a primary reason is correct in the following sense.

Once the expressive resources provided by terms such as 'prefer', 'obliged', and 'ought' are available, it must be possible to use them to make explicit the implicit practical inferential commitment underlying bits of practical reasoning such as  $(\alpha)$ ,  $(\beta)$ , and  $(\gamma)$ . But there is nothing incoherent about an idiom that lacks those expressive resources. Practical reasoning can still take place in it, and there is still a perfectly serviceable distinction between good and bad inferences available within such an idiom. It is by comparing the more primitive practices of giving and asking for reasons for action to the sophisticated ones made possible by the introduction of inference-explicitating locutions such as 'ought' that we can understand (in terms of deontic scorekeeping) the expressive role those locutions play.

The broadly normative or evaluative vocabulary that Davidson understands as expressing the pro-attitudes needed to turn the incomplete reasons offered as premises in  $(\alpha)$ ,  $(\beta)$ , and  $(\gamma)$  into complete reasons is actually used to make explicit in assertible, propositional form the endorsement of a *pat-*



tern of inferences. Different patterns of inference correspond to different sorts of norms or pro-attitudes. Consider someone assessing the propriety of the practical inference in  $\{\alpha\}$ , in a primitive idiom that does not yet permit the formulation of (a). Suppose that the scorekeeper who assesses the practical reasoning attributes to someone commitment to the premise of  $\{\alpha\}$ , and also entitlement to that commitment. The question is whether entitlement to the doxastic commitment serving as the premise is inferentially heritable by the practical commitment serving as the conclusion.

To take it that it is, for a particular interlocutor, just is implicitly to attribute a desire or preference for staying dry. If the inferential commitment that underwrites this piece of practical reasoning is as expressed by (a), then  $\{\alpha\}$  is just one of a whole family of inferences that stand or fall together. For instance, an attributor who takes  $\{\alpha\}$  to be entitlement-preserving will also take the following two inferences and a host of similar ones to have that status.

( $\alpha'$ ) Only standing under the awning will keep me dry,  
so I shall stand under the awning.

( $\alpha''$ ) Only remaining in the car will keep me dry,  
so I shall remain in the car.

To attribute a preference for staying dry to an individual is just to take inferences of this form to be entitlement-preserving, for that individual.

#### 4. Varieties of *Prima Facie* Reasons for Action

Of course there can be competing entitlement-preserving inferences, corresponding to other desires. For recall that permissive inferences generally, whether doxastic or practical, can entitle one to incompatible conclusions<sup>44</sup>—though once an interlocutor endorses one of them, the undertaking of that commitment removes any entitlements that may hitherto have been available for competing claims. So endorsing this pattern of inferences as entitlement-preserving for an individual—which is implicitly attributing the preference that one could explicitly attribute either by attributing commitment to (a) or by undertaking commitment to the ascriptional claim “A desires to stay dry”—does not require attributing to that individual the practical commitment expressed by the conclusion in case commitment to a premise of the proper form is attributed. This is another way of saying that even in the presence of the desire, the belief need not lead to the formation of an intention, for there may be competing desires or other considerations in play. The notion of entitlement-preserving inferences accordingly provides a pragmatic analysis, in deontic scorekeeping terms, of the notion of *prima facie* reasons (whether doxastic or practical).

That a scorekeeper treats inferences of the form common to  $\{\alpha\}$ ,  $\{\alpha'\}$ ,

( $\alpha'$ ) . . . to be entitlement-preserving for interlocutor *A* does not involve any commitment to treating them as entitlement-preserving for interlocutor *B*, even apart from any consideration of the attribution of incompatible commitments (doxastic or practical, or either species of inferential). Treating these inferences as permissively good for *A* but not for *B* is just what attributing the relevant preference to *A* but not to *B* consists in. This is not how endorsement of doxastic inferences (even permissive ones) works. Endorsing a doxastic inference (one whose premises and conclusions are claims, that is, expressions of possible beliefs), treating that inference as entitlement- or commitment-preserving for *one* interlocutor, involves treating it as good for *all* interlocutors—subject, as always, to disqualification by commitment to incompatible claims, and with the proviso that differences in collateral doxastic commitments can make a difference in what premises are available as auxiliary hypotheses in such inferences.

This difference in generality is a fundamental difference between doxastic inference and this sort of practical reasoning. Desire is multifarious and different from individual to individual, but truth is one; so, according to each scorekeeper, the inferences that can be described unofficially as good in the sense of truth-preserving are one, while those practical inferences that are underwritten by desires are many. Of course, beliefs may differ from individual to individual as much as desires, and with it the endorsement of inferences whose propriety is underwritten by particular doxastic commitments, although the social institution of the status of objective information by the interpersonal dimension of assertional authority and the justificatory responsibility to respond to challenges incorporates an implicit norm of common belief that has no analog for desire. The difference being pointed to here is rather that attributions of *conative* commitments are construed here as fundamentally a kind of inferential commitment, linking doxastic and practical commitments, while *cognitive* or doxastic commitments and practical commitments are inferentially articulated and inferentially significant, but not themselves inferential commitments. That (in informal terms) desires vary from individual to individual, as beliefs do, is accordingly reflected in a structure of inferential commitments in the conative case that differs from that of the cognitive case.

Permissive proprieties of practical reasoning, endorsement of which is implicitly attributing—or in the reasoner's own case, acknowledging (which is self-attributing)—desires or preferences (pro-attitudes in a strict sense, as represented by the example of ( $\alpha$ )), are, however, only one species. Those represented by the example of ( $\beta$ ) need not be understood as having anything in particular to do with desires or preferences. The norm, rule, or requirement that bank employees wear neckties is what makes going to work into a reason for wearing a necktie, for bank employees. Taking it that there is such a norm or requirement just is endorsing a pattern of practical reasoning—namely, taking ( $\beta$ ) to be an entitlement-preserving inference for anyone

who is a bank employee. This inferential pattern is different from that exhibited by  $(\alpha)$  in two ways. First, there need not be for each interlocutor for whom  $(\beta)$  is taken to be a good inference a set of other inferences corresponding to  $(\alpha)$ ,  $(\alpha')$ ,  $(\alpha'')$ . Second, the scorekeeper will take  $(\beta)$  to be a good inference for any interlocutor  $A$  such that the scorekeeper *undertakes* doxastic commitment to the claim that  $A$  is a bank employee. Thus the way in which the scorekeeper's endorsement of the inference  $(\beta)$  as permissively good is distributed across various interlocutors who might reason that this way is different from that of  $(\alpha)$ .

Here the norm implicitly underwriting the inference is associated with having a certain status, as employee of a bank, rather than with exhibiting a certain desire or preference. Whether one has a good reason to wear a necktie just depends on whether or not one occupies the status in question. This pattern—where what matters is the scorekeeper's undertaking of a commitment to  $A$ 's occupying the status, rather than  $A$ 's acknowledgment of that commitment—corresponds to an objective sense of 'good reason for action' (according to the scorekeeper). In this sense, that  $A$  is preparing to go to work can be a good reason for  $A$  to wear a necktie, even though  $A$  is not in a position to appreciate it as such. The scorekeeper might take it that  $A$  is entitled to a practical commitment to wear a necktie, even though  $A$  could not justify it by producing the reasoning in  $(\beta)$ . In wearing a necktie,  $A$  would be acting *with* a reason, even if not *for* a reason. This corresponds to taking a reliable noninferential reporter to be entitled by that reliability to various observations, even in the case where the reporter is not in a position to appeal to that reliability in justifying those claims.

In another sense, of course, for the norm that the scorekeeper takes to be in force to supply a reason *for*  $A$ , the claim that  $A$  is a bank employee must also be acknowledged by  $A$ . For  $A$  to be able to justify a commitment to wearing a necktie by rehearsing the reasoning of  $(\beta)$ ,  $A$  must also endorse the pattern of inference codified explicitly in  $(b)$ . For a scorekeeper to take  $A$  to have a good practical reason in this stronger sense—that not just the scorekeeper, but  $A$  could produce it—requires that the scorekeeper attribute to  $A$  endorsement of an *inference*. In the model as presented thus far, this can be done only by attributing commitment to a claim codifying that inference. Depending on the expressive resources available, this might either be  $(b)$  or a set of corresponding conditionals.

One final stronger sense of *reason for*  $A$  is sometimes invoked by philosophers who insist that even  $(b)$  together with  $A$ 's acknowledgment of being a bank employee fall short of providing one unless supplemented by  $A$ 's desire to do what is required as a bank employee. It is indeed always possible, by supplying "suppressed" premises as needed, to assimilate all practical reasonings to the form of  $(\alpha) + (a)$  (assimilating them to belief-desire reasons), so that norms and evaluations appear only in the role of objects of preference, as *staying dry* does in  $(\alpha) + (a)$ . And it remains true that the role of  $(\beta)$  would

be quite different if (b) were conjoined with some claim *incompatible* with attributing to *A* such a desire. But this is just another instance of the phenomenon mentioned above as motivating but not warranting an enthymematic view. The inference from *p* to *q* is also undercut by the *denial* of the conditional  $p \rightarrow q$ , but this does not show that the conditional is a suppressed premise in the original argument. The present approach requires no such instrumental reductionism, however, for it is possible to say what it is for scorekeepers to *treat* various other (from this point of view truncated) forms of practical reasoning *as* correct, as entitling agents to their practical commitments. Doing so makes it possible to explain how various normative vocabulary works (what it expresses)—including the vocabulary of commitment and entitlement that is employed in laying out the deontic scorekeeping model of discursive practice.

Endorsement of practical reasoning of the sort of which ( $\gamma$ ) is representative, codified in the form of a normative principle by (c), corresponds to an inferential commitment exhibiting a pattern different from those involved in either ( $\alpha$ ) or ( $\beta$ ). For a scorekeeper who takes ( $\gamma$ ) to be entitlement-preserving for *A* takes it to be entitlement-preserving for *anyone*—regardless of desires or preferences and regardless of social status. Inferential commitments displaying this pattern are made explicit by *unconditional* 'ought's, whereas those displaying the other two patterns are made explicit by *prudential* 'ought's (in the case of (a)) and *institutional* 'ought's (in the case of (b)). Unconditional 'ought's, which correspond to this agent- and status-blind pattern of endorsement of practical inferences as entitlement-preserving, are one candidate that has been proposed as a good thing to mean by "moral 'ought'." Some thinkers insist rather that to treat reasons as moral reasons requires treating them as overriding; this amounts to saying that the 'ought' in (c) is a moral 'ought' only if ( $\gamma$ ) is not only entitlement-preserving but also *commitment-preserving*—that anyone committed to the doxastic premises is thereby committed to the practical conclusion. It is not the point of this discussion to take a stand on how to distinguish specifically moral norms. Nor is the point to try to provide an exhaustive catalog of the sorts of norms (or "pro-attitudes," in the broad sense) there can be. The point is just to show that various important sorts of norms (or pro-attitudes) can sensibly be thought of in deontic scorekeeping terms as corresponding to different patterns of endorsement of practical *inferences*.

To endorse a practical inference as entitlement-preserving is to take the doxastic premises as providing reasons for the practical conclusion. To exhibit a piece of good practical reasoning whose conclusion is a certain intention is to exhibit that intention, and the action (if any) that it elicits, as *rational*—that is, as reasonable in the light of the facts cited and the commitments exhibited in the premises. So *all* of the 'ought's that make explicit species of practical reasoning taken as examples here (the prudential or preferential 'ought', the social or institutional 'ought', and the unconditional

'ought') are different kinds of *rational* 'ought'. Being rational is just being in the space of giving and asking for reasons, and being a rational *agent* is being in the space of giving and asking for reasons for what one *does*. When the proprieties of practical inference that articulate that space are made explicit in the form of claims, they take the form of norms—of rational 'ought's. Rationality is the genus to which all these species of 'ought's belong.

There is no a priori reason to identify the rational with some one of the species of practical reasoning (for instance the prudential) and cut and paste the rest into suitable shape to be assimilated to the favored one. Being a reason is to be understood in the first instance in terms of what it is for a community to *treat* something in practice as such a reason, on the practical side of reasons for action just as on the doxastic side of reasons for claims. In neither case is this approach to normative status (what one is *really* entitled or committed to) through normative attitude (what one is *taken* to be entitled or committed to) incompatible with making eventual sense of *objective* norms, which underwrite the possibility that everyone's attitudes toward them are wrong. But understanding what is meant by such objective proprieties—what is *really* a good reason, as opposed to just what is *treated* as one—comes at the end of the story.<sup>45</sup> It is not something that can be understood a priori and imposed as a constraint at the outset.

## VI. INTENTIONS

### 1. *Reasons and Entitlement to Practical Commitments*

Exhibiting a piece of practical reasoning *rationalizes* the practical commitment or intention that is its conclusion. It displays reasons for that intention, offers a rational justification for it, shows how one might become rationally entitled to it. Accepting a practical inference as entitling someone to a practical commitment in this sense requires endorsing the *inference* as permissively good (and so only as providing a *prima facie* case for commitment to the conclusion, defeasible by incompatible commitments) for the agent whose conduct is being assessed. It does not require that the inference be accepted as one that would be (permissively) good in the scorekeeper's own case; the scorekeeper need not share the desire, preference, or institutional status that is implicitly attributed by treating some practical inferences as good for some agents. Nor does it require that in all cases the scorekeeper assessing that entitlement endorse the *premises*; a requirement of that sort picks out the special sense of objective entitlement. As Davidson says about prudential or preferential practical reasoning (the only kind he acknowledges): "When we talk of reasons in this way, we do not require that the reasons be good ones. We learn something about a man's reasons for starting a war when we learn that he did it with the intention of ending all wars [for Davidson this is equivalent to 'because he desired to end all wars'],

even if we know that his belief that starting a war would end all wars was false. Similarly, a desire to humiliate an acquaintance may be someone's reason for cutting him at a party though an observer might, in a more normative vein, think that that was no reason. The falsity of a belief, or the patent wrongness of a value or desire, does not disqualify the belief or desire from providing an explanatory reason."<sup>46</sup>

An agent can fulfill the justificatory responsibility involved in undertaking a practical commitment by exhibiting a piece of practical reasoning in which a commitment with that content serves as the conclusion. In the doxastic case, what is an entitling justification for one is an entitling justification for all, except for disqualifications due to commitment to claims incompatible with the premises or the conclusion. In the practical case, entitling justifications need not be portable across agents in this way. Displaying an intention as the conclusion of a piece of practical reasoning that is good in this sense makes it intelligible by showing reasons that could entitle the agent to it. To secure an attribution of entitlement to a practical commitment in this way, the practical reasoning in question may be offered as a justification by the agent, perhaps upon being challenged to do so. But it also may be attributed by the scorekeeper, who constructs the practical argument from premises already attributed to that agent, according to patterns of practical inference the scorekeeper endorses for that agent. These might be patterns of inference the scorekeeper endorses for everyone (unconditional 'ought's), or endorses for the agent on the basis of status (institutional 'ought's), or endorses only for the agent, thereby implicitly attributing idiosyncratic desires or preferences (prudential 'ought's), or of some other kind.

A scorekeeper who in this way takes an agent to be entitled to a practical commitment on the basis of its being the conclusion of a practical inference taken to be good for the agent, and who also attributes to that agent commitment to the premises of that inference, need not take it that acknowledgment of the practical commitment actually arose as the result of a process of inference by the agent from acknowledgment of those premises. As Davidson says "We cannot suppose that whenever an agent acts intentionally, he goes through a process of deliberation or reasoning, marshals evidence, and draws conclusions. Nevertheless, if someone acts with an intention, he must have attitudes and beliefs from which, had he been aware of them at the time, he *could* have reasoned that his action was desirable (or had some other positive attribute)."<sup>47</sup> For Davidson, acting intentionally and acting for reasons are the same thing. From the present point of view, this position involves conflating the two deontic statuses of practical commitment and entitlement to such a commitment. An act is intentional if it is (or is, as the exercise of a reliable differential responsive disposition, noninferentially elicited by) the acknowledgment of a practical commitment. To act for reasons is to be *entitled* to that practical commitment. One can in particular cases act intentionally but without reasons, even though there is no making sense of inten-

tions apart from their liability to the demand for reasons. But Davidson's point survives this confusion. An intention can be rendered intelligible as rational (a practical commitment can be displayed as one the agent is entitled to) by displaying reasons for it (premises from which it *could* legitimately have resulted as the conclusion of a good practical inference), even in cases where it was not in fact arrived at by such a process.

Contemporary thought about action begins with Anscombe's insight, developed with great force and clarity by Davidson, that the difference between actions and other performances—the answer to Wittgenstein's challenge to explain the difference between my raising my arm and my arm going up—is that actions are performances that are intentional under some description.<sup>48</sup> Any performance can be specified in many ways: Davidson moves his finger, flips the switch, turns on the light, alerts the burglar, causes a short-circuit in the wiring, starts a fire. These are all things he *does*; the different descriptions are different ways of specifying one action he performs. Not all of these are specifications under which what he does is intentional. But they are all specifications of an action, so long as what he does is intentional under *some* specification, for instance turning on the lights. In Davidson's slogan, being an action is an *extensional* property of an event (whether a given event is an action or not is not sensitive to how the event is specified), while being intentional is an *intensional* property of an event (whether a given event is intentional or not is sensitive to how the event is described). The very same event is intentional as turning on the lights but unintentional as alerting the burglar, causing a short-circuit, and starting a fire. The extensional property of performances, *being an action*, is defined in terms of the intensional property of performances of *being intentional* by existential quantification over descriptions or specifications of the performance: if it is intentional under any one of them, it is an action under all of them.

By this strategy the problem of explaining what privileges some (but not all) of an agent's performances as actions is reduced to the problem of explaining what privileges some (but not all) descriptions or specifications of an action as ones under which a performance is intentional. Davidson's solution to this problem in turn is that a performance is intentional under a description if that description figures as the conclusion of a piece of practical reasoning that exhibits the agent's reasons for producing it. These two moves together—the account of actions as performances that are intentional under some description, and the account of performances as intentional under just the descriptions that appear as the conclusions of practical inferences that rationalize those performances by giving reasons for them—reduce the problem of explaining what is special about action to that of explaining the giving of practical reasons. Davidson's account of primary reasons as pairs of beliefs and pro-attitudes is then offered to explain what it is for reasons to rationalize a performance according to a practical inference.

The Davidsonian explanatory structure provides a recipe, then, for turning

an account of practical reasoning into an account of action. In such a context, the present account of practical reasoning in terms of deontic scorekeeping on inferentially articulated practical commitments and entitlements to such commitments has some advantages over the one Davidson himself endorses. As originally presented, Davidson's theory eschews intentions entirely, in favor of beliefs and desires: "The expression 'the intention with which James went to church' has the outward form of a description, but in fact it is syncategorematic and cannot be taken to refer to an entity, state, disposition, or event. Its function in this context is to generate new descriptions of actions in terms of reasons; thus 'James went to church with the intention of pleasing his mother' yields a new, and fuller, description of the action described in 'James went to church'."<sup>49</sup> The account offered here, by contrast, explains deontic statuses corresponding to beliefs and intentions and defines those corresponding to desires (and other <sup>s</sup>expressions of pro-attitudes<sup>s</sup>) in terms of them.

Acknowledging intendings as full-fledged intentional states (or attitudes toward deontic statuses) avoids at the outset a difficulty that forced Davidson to modify his earlier account. For there are cases where someone has an intention (not just a reason for action), but no action arises from it. The possibility of intention without action is a symptom of the limited scope of Davidson's original discussion. As he puts the point in his introduction to the collection of his essays on this topic: "When I wrote ['Actions, Reasons, and Causes'] I believed that of the three main uses of the concept of intention distinguished by Anscombe (acting with an intention, acting intentionally, and intending to act), the first was most basic. Acting intentionally, I argued . . . was just acting with some intention. That left intending, which I somehow thought would be simple to understand in terms of the others. I was wrong . . . Contrary to my original view, it came to seem the basic notion on which the others depend; and what progress I made on it partially undermined an important theme of ['Actions, Reasons, and Causes']—that 'the intention with which the action was done' does not refer to an entity or state of any kind."<sup>50</sup>

## 2. Two Sorts of Intention

Explaining intentional action requires only what Searle calls *intentions in action*. Explaining pure intending requires also what he calls *prior intentions*. The distinction is motivated by the fact that "I can do something intentionally without having formed a prior intention to do it, and I can have a prior intention to do something and yet not act on that intention."<sup>51</sup> Pure intendings are special cases of prior intentions. In the deontic idiom, both sorts of intentions are (acknowledgments of) practical commitments—that is, commitments to act. Cases of intentions in action without prior intentions are those in which the performance that is accorded by a scorekeeper



the significance of an acknowledgment of a practical commitment is just the action itself. Consider Searle's example: "Suppose I am sitting in a chair reflecting on a philosophical problem, and I suddenly get up and start pacing about the room. My getting up and pacing about are clearly intentional actions, but in order to do them I do not need to form an intention to do them prior to doing them. I don't in any sense have to have a plan to get up and pace about. Like many of the things one does, I just do these actions; I just act."

In scorekeeping terms, undertaking a commitment is doing anything that makes it appropriate for it to be attributed. This may involve a distinct explicit acknowledgment, as in the case of asserting; it may be consequential, as in undertaking commitment to the consequences of a claim that is asserted; or it may be a default matter. Scorekeepers will take one not only to be entitled but to be committed to the claims that orange is a color and that there have been black dogs, unless these have been overtly disavowed (and even this may not be sufficient, if collateral commitments that are not disavowed entail these claims). The sort of case Searle considers concerns actions that have the default significance of intentional actions, as ones accompanied by a commitment for which the question of entitlement by reasons is in principle in order. It might turn out that what is, when described in a suitably impoverished vocabulary of motions of limbs, exactly similar behavior is not intentional at all but automatic, involuntary, compulsive behavior, triggered ineluctably by pheromones. In that case the default attribution of a practical commitment would be defeated. But the undertaking or acknowledging of a commitment to act need not be a performance distinguishable from the act one is committed to perform.

Nonetheless, in many cases, it is a separately datable event. In such cases of prior intention, acknowledging the commitment antedates fulfilling it (or not, as the case may be). I can *now* acknowledge a commitment to get on the bus *when* it arrives. My mastery of the practical conceptual (because inferentially articulated) content of that commitment includes my mastery of the *noninferential* differential disposition to respond to it and the joint fulfillment of its condition (the bus arriving) by getting on the bus. When the bus arrives, the condition is fulfilled. My getting on the bus is an acknowledgment of a practical commitment to get on the bus *now*—an intention in action. If I have the reliable noninferential differential disposition to respond to the acknowledgment of a practical commitment to do *A* when (if) *C* by doing *A* when (if) *C*, then my prior intention to get on the bus when it arrives will *mature* into a corresponding intention in action (marked above by the 'now' in the linguistic expression of the intention).

Prior intentions must specify the actions one is committing oneself to perform in *general* terms; they would be expressed by statements of the form

I shall get on the bus when it arrives,

in which the action is specified by a description. Intentions in action can be thought of as specifying the action one is committing oneself to perform in *demonstrative* terms; they could be expressed by statements of the form

I shall *now* do *this*.

Only an intention in action can be directed at a particular unrepeatable action. The process by which a prior intention ripens into an intention in action—the exercise of the practical skill of reliably responding to the undertaking of a commitment to bring about a state of affairs of such and such a description when such and such conditions obtain by doing so—involves recognizing when doing *this now* would fulfill such a commitment. (See further discussion below in 8.5.2.) What are here called “prior intentions,” Sellars calls simply “intentions,” and what are here called “intentions in action,” Sellars calls “volitions”:

A simple case of the relation of intending to volition can be illustrated by considering Jones, who has formed the intention of raising his hand in ten minutes. Suppose that no alternative course of action recommends itself to him. Then we may picture the situation as follows:

I shall raise my hand in *ten* minutes.

. . .

I shall raise my hand in *nine* minutes.

. . .

I shall raise my hand *now*.

(which culminates in action, if Jones happens not to be paralyzed).<sup>52</sup>

In Sellars’s idiom, a volition is an intention whose time has come.

Sellars takes the capacity respond reliably to prior intentions whose time has come by the formation of intentions in action (here, acknowledgments in the form of suitable performances) to be part of grasping the meaning of what is expressed by ‘shall’, and of the practical content of the particular concepts that articulate the content of those intentions. As in the account endorsed here, he takes these capacities to be part of the “important similarity between learning to make the language-entry transition of responding to presented red objects by saying ‘This is red,’ and learning the language-departure transition” involved in exercising those capacities.<sup>53</sup> Just as in the case of language entries or noninferentially elicited but inferentially articulated doxastic commitments, the existence of reliable differential responsive dispositions is compatible with making mistakes, so in the case of language exits or performances noninferentially elicited by inferentially articulated practical commitments, the existence of reliable differential responsive dispositions is compatible with failure. Mistakes of observation are diagnosed by scorekeepers by comparison of the contents of the doxastic commitments *attributed* to the observer and those *undertaken* (whether noninferentially,

inferentially, or by testimony) by the scorekeeper—for instance when the scorekeeper attributes a commitment to the claim that the ball is red and undertakes a commitment incompatible with the claim that the ball is red. Failures of action are diagnosed by scorekeepers by comparison of the contents of the practical commitments *attributed* to the agent and the contents of the doxastic commitments *undertaken* by the scorekeeper.

For instance the scorekeeper who attributes a commitment that could be expressed (Sellars-wise) as

Shall [The basketball goes through the hoop]

may also be obliged (perhaps observationally) to acknowledge a separate commitment incompatible with the claim that the basketball goes through the hoop (for example that expressed by “The ball missed the hoop”). A practical commitment may also remain a “pure” intending, eventuating in no action, successful or unsuccessful, for one of two reasons. It might be that the condition of maturation of a prior intention into an intention in action never is satisfied—the bus never arrives, the commitment expires. Or it may be that the commitment is relinquished, perhaps in favor of one with an incompatible content—the agent undergoes a change of mind.

### 3. *Acknowledgments of Commitments Can Cause and Be Caused*

When prior intentions are made explicit, they specify in general descriptive terms the performance the agent is committed to produce. When intentions in action, which are implicit in suitable performances, are made explicit, they specify those performances demonstratively. The way in which prior intentions elicit suitable performances, and so intentions in action, is a causal process. The analogy between action and observation (between language-exit transitions and language-entry transitions) is intended to illuminate the nature of the process involved.

Davidson's original essay famously endorses the claim that reasons are causes. The difference between a commitment's being *a* reason for an agent's action and its being *the* reason for that action must be explained in terms of differences in the causal roles played by various states. Primary reasons, conceived of as pairs of a set of beliefs and a set of pro-attitudes, rationalize actions (which accordingly count as intentional) first by providing reasons for them and second by serving to bring them about (“in the right way”). The account offered there has been criticized here for running together the notion of being *committed* to act in a certain way and being *entitled* to do so by reasons; even though the first deontic status cannot be made sense of apart from the second (any more than doxastic commitments can be made sense of apart from practices of giving and asking for reasons entitling interlocutors

to them), nonetheless these are distinct statuses, and one can be committed without being entitled, can act intentionally without having reasons for doing so. This distinction between practical commitments and entitlements to such commitments also, it has been claimed, opens up a space for the notion of *prior* intention (besides that of intention in action), of which pure or unconsummated intendings are a species.

What becomes of the doctrine that reasons are causes, when intentional states are construed in terms of social scorekeeping on deontic statuses and the deontic statuses of commitment and entitlement are appropriately distinguished? The account of the social practices that institute deontic statuses appeals to such statuses only as scorekeeping devices. The significance of being committed to a certain claim or assertible content is normative. It has to do with what else one is committed or entitled to. It is articulated by *proprieties* of scorekeeping and consists of the *proper* antecedents and consequences of that status. In the same way, that there are two strikes on a batter is a status properly acquired by various performances (just which depending on the antecedent score), a status that alters the significance various further performances have for the subsequent score.

Any effect that such elements of the score have on what performances are *actually* produced is indirect, mediated by the *attitudes* of those who keep score. The score determines only what *ought* to be done, what would be *proper*. What ought to be done and what is proper affect what players do only insofar as they are trained to respond in various ways to *taking* a certain course of action to be proper. The only access that deontic statuses have to the causal order is through the deontic attitudes of the scorekeeping practitioners.

Inferential relations among propositional contents are a matter of normative relations among deontic statuses: commitment to the claim that lions are mammals entails commitment to the claim that lions are vertebrates. Inferring, by contrast, is a causal process that relates deontic *attitudes: acknowledging* (and equally, attributing to someone else) commitment to the claim that lions are mammals will, under various circumstances and in those well versed practically in the inferential relations among deontic statuses, have as a causal consequence acknowledging (or, correspondingly, attributing) commitment to the claim that lions are vertebrates. Unless the members of a linguistic community are pretty good at keeping score by altering their attitudes as they ought to according to the contents associated with the deontic statuses in terms of which they keep score, there is no point in interpreting them as engaging in the practices specified by those proprieties of scorekeeping. Nonetheless, normative status is one thing, the attitudes of attributing and undertaking those statuses, the alteration of which is what scorekeeping consists in, is another.

As it is with the inferential articulation of the conceptual contents conferred on states, attitudes, performances, and expressions by deontic score-

keeping practices, so it is with their noninferential involvements, which confer empirical and practical conceptual contents on them. What observable states of affairs causally elicit in perception, according to reliable differential responsive dispositions, is in the first instance deontic *attitudes* rather than statuses: *acknowledgments* of doxastic commitments. What in action causally elicits the production of performable states of affairs (by the exercise of reliable differential responsive dispositions) is in the first instance deontic *attitudes* rather than statuses: *acknowledgments* of practical commitments. ("In the first instance" because acknowledging a commitment is one way of undertaking one, so those deontic attitudes have scorekeeping consequences for the deontic statuses of those whose attitudes they are.) That a particular doxastic commitment was elicited by the exercise of such dispositions is another way of putting the condition on perception that the belief not only be caused by the state of affairs reported but be caused by it "in the right way." That a particular practical commitment elicits a performance by the exercise of such dispositions is another way of putting the condition on action that the performance not only be caused by the intention but be caused by it "in the right way." Mastering the two sorts of reliable differential responsive dispositions connecting noninferentially acquired acknowledgments of doxastic commitments to their appropriate causal antecedents and noninferentially efficacious acknowledgments of practical commitments to their appropriate causal consequents is part of grasping, in one's scorekeeping practice, the empirical and practical components of the contents of concepts employed in observation and action, and of those theoretical concepts inferentially related to these.

Thus just as 'belief' is ambiguous in scorekeeping terms, referring sometimes to a deontic status and sometimes to a deontic attitude (sometimes to doxastic commitment and sometimes to acknowledgment of such a commitment), so 'intention' is ambiguous in scorekeeping terms, referring sometimes to a deontic status and sometimes to a deontic attitude (sometimes to practical commitment and sometimes to acknowledgment of such a commitment). Believing in the sense that entails one's readiness to avow what one believes and to act on it corresponds to acknowledging a doxastic commitment. Intending in the sense that entails one's readiness to act on it (and, should the expressive resources for doing so exist in the linguistic practices in question, to avow it with a 'shall' claim) corresponds to acknowledging a practical commitment. In this sense of 'intention', then, intentions are causes, for in the properly trained agent, acknowledgments of practical commitments reliably causally elicit performances. In this sense of 'belief', when beliefs provide reasons that entitle one to a practical commitment, they *may* function also as causes. They do just in case the acknowledgment of the practical commitment in fact arose by inferring it from an acknowledgment of the belief playing the role of premise in practical reasoning.

One, however, may have intentions without reasons, practical commit-

ments to which the agent is not entitled by doxastic commitments suitably related as premises of practical inferences. In that case there can still be action, but it will not be caused by reasons. Again, the agent may be entitled to the practical commitment, according to a scorekeeper, by doxastic commitments that are attributed by the scorekeeper but not acknowledged by the agent. This might happen where the scorekeeper takes the agent to have undertaken corresponding doxastic commitments as inferential consequences of others that are acknowledged, but where the agent has never been through the process of inference that would lead to acknowledging those consequences. Commitments of this sort could still entitle the agent to the practical commitment, even though only the scorekeeper, and not the agent, would be in a position to exhibit the practical reasoning that secures that entitlement. In such cases, too, the agent's reasons for the action would not be functioning as causes. So once the deontic statuses of commitment and entitlement to commitments are properly sorted out, it turns out that a nonlinguistic performance can have at least three different sorts of scorekeeping significance:

1. an agent's acting intentionally—that is, acknowledging a practical commitment by producing a performance or exercising a reliable non-inferential differential disposition to respond to acknowledgments of practical commitments by producing a performance,
2. an agent's having reasons for action or acting *with* reasons—that is being entitled to a practical commitment, and
3. an agent's acting *for* reasons, the action being caused by (attitudes toward) reasons for action—that is the acknowledgment of the practical commitment having arisen by a process of inference from acknowledgment of the commitments that provide the entitling reasons.

The first does not entail the second (nor vice versa), nor does the second entail the third, though they are all compatible; one can act intentionally either with or without reasons, and one may or may not act for the reasons one has.

#### *4. Acknowledging Commitments Need Not Be Modeled on Promising*

Davidson's own view about intentions (once he comes to countenance them at all) identifies them as all-things-considered judgments, in the light of all the agent's primary-reason-providing beliefs and desires, that an action of a certain kind is desirable, good, or ought to be performed.<sup>54</sup> From the present point of view this is an unsatisfactory conclusion, both because of its appeal to unanalyzed notions of desirability, good, and what ought to be done, and because it does not say what it is for the attitudes these locutions express to become explicit in the form of an evaluative judgment.

He has not laid out the practices that could confer such a content on a judgment, so he has not explained how to understand fully the implicit commitment that is being made propositionally explicit by the use of this vocabulary. Such a complaint is of exactly as much interest as the concrete alternative account that is recommended; it is better to light a candle than to curse the darkness. The deontic scorekeeping account of practical commitments, built on that of doxastic commitments, is meant to supply this want. In order to recommend that idiom over the one Davidson endorses, however, it is necessary to confront the argument against construing intentions as a sort of commitment that he offers along the way to his identification of intentions as a special kind of judgment.

The leading idea of the present account is that acting intentionally is doing something that has the deontic scorekeeping significance of acknowledging a practical commitment (in the case of intentions in action), or noninferentially producing a performance by exercising a reliable differential disposition to respond to the acknowledgment of a practical commitment (where prior intentions are involved). Intentions are identified with such acknowledgments of commitments, and the reasons for or with which an agent acts with the attitudes or facts that entitle that agent to those practical commitments, according to the role they play as premises in practical inferences. This normative (more specifically deontic) approach to intention and action is rooted in Sellars's discussion of the giving and asking for reasons for action, which has been elaborated along different lines by Castañeda.<sup>55</sup> Although the details of their accounts are different, the overall approach is very similar. Sellars never actually talks about intentions in terms of commitments, but this way of putting it is implicit in his account.<sup>56</sup> If there is something wrong with thinking about intentions in terms of commitments, then this whole approach is broken-backed. So it is of the first importance to consider Davidson's arguments against it.

Davidson begins by considering theories that focus on the speech act of expressing an intention (the speech act that Sellars regiments using 'shall'). He observes that "saying, under appropriate circumstances, that one intends to do something, or that one will do it, can commit one to doing it; if the deed does not follow, it is appropriate to ask for an explanation."<sup>57</sup> The suggestion that forming an intention is performing a speech act of this sort (perhaps addressed to oneself)—a performative theory of intention—is rejected because "the performative character of commands and promises which makes certain speech acts surprisingly momentous depends on highly specific conventions, and there are no such conventions governing the formation of intentions."<sup>58</sup> Indeed it seems enough to observe that, although for Davidson as for the deontic scorekeeping account, one must be able to talk in order to have intentions (because it is only in the context of linguistic practices of giving and asking for reasons that anything could be accorded the significance of an intention), there is no necessity that there actually be a term 'shall' that

overtly marks a special speech act that attaches to propositional contents the significance of explicit undertakings of practical commitments. The present explanatory strategy demands rather that that implicit force first be explained in scorekeeping terms so that a clear sense can be made of the introduction of locutions whose expressive function is to make that force explicit in the form of an assertion. The present account does not understand the undertaking of a practical commitment as requiring a special speech act. What is important is the attitude of acknowledging a commitment to act; any connection with special performative speech acts comes later.

But Davidson objects as well to the invocation of commitments in this context: "Promising involves assuming an obligation, but even if there are obligations to oneself, intending does not normally create one. If an agent does not do what he intended to do, he does not normally owe himself an explanation or apology, especially if he just changed his mind; yet this is just the case that calls for explanation or apology when a promise has been broken. A command may be disobeyed, but only while it is in force. But if he does not do what he intended because he has changed his mind, the original intention is no longer in force."<sup>59</sup> There are a number of points being made here; they turn on disanalogies between forming an intention and making a promise, which serves for Davidson as the paradigm of the undertaking of a commitment. To begin with, he offers the implicit suggestion that there may be problems with the notion of making a promise to oneself. Promises (like commands) are made *to* someone, while no one else is typically addressed by the formation of an intention. Then the central objection is presented, that there seems to be no *sanction* associated with failure to perform as one is committed to perform. If there is a commitment, then fulfilling it or failing to fulfill it ought to make some sort of difference, as it does in the case of failing to fulfill a promise or to carry out an order from a suitable authority. Yet once an intention has been formed, it can be withdrawn without penalty—the agent can have a change of mind. Promises would not be promises, would not involve the undertaking of *commitments*, if they could be canceled at the whim of the promiser. How could sense be made of a commitment that was in force only as long as the one committed decided to keep it in force but that could be relinquished without penalty at any time? Davidson concludes that the disanalogies are too great and that forming an intention cannot sensibly be conceived as undertaking a commitment.

These disanalogies between intending and promising, even promising oneself, should be acknowledged. But the conclusion follows only if there is no other model of acknowledging or undertaking commitments available besides that of promising. The deontic scorekeeping account of acknowledging *assertional* or *doxastic* commitments shows that this is far from being the case. Assertional commitments, after all, can be withdrawn without penalty by the asserter who undergoes a change of mind. Commitments of this sort



are put in force by the performance of a speech act and, except for special cases (where one is also consequentially committed to the claim in virtue of other, unwithdrawn assertions), can be relinquished without penalty by another speech act. It is true that nothing resembling promising could work this way, but the model of asserting shows that there are other ways to conceive the undertaking and acknowledging of commitments. While a doxastic commitment is in force, that fact has consequences; the undertaking of such a commitment has a significance for the deontic score. Commitment to one content entails commitment to others and precludes entitlement to yet others.

It is the same with practical commitments as here presented. Undertaking one is not without significance simply because it can be voided, withdrawn, or overridden essentially at the whim of the agent. For when such a commitment is in force (according to a scorekeeper who attributes it), it is significant. It entails various further commitments and precludes various entitlements. It can license the attribution of doxastic commitments (standing in for beliefs) that would warrant it, according to an attributed piece of practical reasoning. Like doxastic commitments, practical commitments involve a (conditional) justificatory responsibility to vindicate the commitment by demonstrating entitlement to it (upon suitable challenge). This forms part of the significance of these commitments, on the side of antecedents rather than of consequents, for it determines the circumstances under which it is appropriate to acquire these statuses. The disanalogies between promising and undertaking doxastic commitments do not make the latter sort of status unintelligible as a species of commitment, and the same disanalogies between promising and the undertaking of practical commitments, which are modeled closely on doxastic ones, do not make that sort of status unintelligible as a species of commitment.

It might be objected that the disanalogies between doxastic and practical commitment reinstate the difficulty. For on the one hand, doxastic commitments are like those undertaken by promising, and unlike intendings, in that they are intelligible only in terms of a speech act that has the significance of an overt public acknowledgment of them. And on the other hand, a scorekeeping sanction for failing to fulfill the justificatory responsibility associated with undertaking a doxastic commitment is the loss (in the eyes of the scorekeeper who attributes the failure of entitlement) of its authority, its capacity to license commitment by others to that same content. But the lack of this sanction is precisely one of the important points of disanalogy between doxastic and practical commitments.

Each of these points might have force if practical commitments were conceived as autonomous—that is, as statuses that could be instituted by practices that did not also institute doxastic deontic statuses. This sort of autonomy is claimed only for doxastic commitments and entitlements, however, not for their practical counterparts. In this context, the disanalogies

between doxastic and practical deontic statuses do not reinstate Davidson's objections. The first point just shows that because of the essential role played by the overt public acknowledgment of doxastic commitments by the performance of speech acts accorded the significance of assertions, doxastic commitments are more like those undertaken by promising than are practical commitments. For practical commitments need stand in no such intimate relation to speech acts accorded the significance of acknowledgments of them. Explicitating locutions permitting the production of speech acts of this sort can be introduced, but practical commitments are intelligible even in their absence. Practical commitments as here conceived are unintelligible apart from all reference to the overt undertaking of commitments by speech acts; that is why they are an essentially linguistic phenomenon. But as here described, the only sort of speech act they presuppose is assertion, the acknowledgment not of practical but of doxastic commitments.

The second point was that intentions do not implicitly claim the sort of interpersonal authority that assertions do, so that the consequences of attributing commitment without entitlement cannot be in the practical case what they are in the doxastic case, namely the undercutting of that authority. But this does not show that entitlement to practical commitments is unintelligible, only that it is different in detail from entitlement to doxastic commitments. Entitlement to practical commitments still has an intrapersonal significance, for instance in connection with the incompatibility of practical commitments (which is linked to that of doxastic commitments), and so with their permissive entailments. Again, this feature does not threaten to make practical commitments unintelligibly private, both because of the irreducibly social character of the deontic scorekeeping, in terms of which such statuses are explained by theory and sustained by practice, and because of the connection with doxastic commitments via practical reasoning.

Entitlements aside, treating a performance as (or as elicited by) an acknowledgment of a practical commitment—that is, treating it as intentional—has scorekeeping consequences not only for the attribution of further practical commitments but also for the attribution of doxastic commitments. It is often possible to infer what an agent believes from what that agent does. Committing oneself to act in a certain way may be committing oneself to taking the world to be a certain way, in the eyes of a scorekeeper who attributes a suitable background of other commitments (of both discursive species). The doxastic commitments an agent is taken consequentially to have undertaken in this way may be incompatible with other doxastic commitments the scorekeeper attributes, in which case entitlement to all of them is undercut. So in part in virtue of the intimate connections between them, the asymmetries between practical and doxastic commitments do not threaten the intelligibility of the deontic scorekeeping significance of the former.

### 5. 'Should' and 'Shall'

Once these points are appreciated, it becomes clear that Davidson's considerations provide no reason not to understand forming an intention as acknowledging a commitment, provided commitment is properly understood according to the practical deontic scorekeeping model of doxastic commitment rather than on the model of promising. But it also becomes clear that there is surprisingly little difference between his construal of intentions as all-out evaluative judgments and the deontic scorekeeping construal of them as acknowledgments of practical commitments. For, first, though Davidson does not think of them that way, in the context of the model presented here, taking intentions to be a kind of judgment *is* taking them to be a kind of commitment. Davidson uses 'judgment' as the genus and allows cognitive and conative, or descriptive and evaluative, species corresponding to beliefs and intentions. The idiom presented here uses 'commitment' as the genus and allows doxastic and practical species corresponding to beliefs and intentions.

For Davidson, there are two sorts of "evaluative" judgment: those that express pro-attitudes suited to be elements of primary reasons for action (that is, those that express merely *prima facie* or *ceteris paribus* evaluations), and those that express intentions and are directly responded to by the production of suitable performances. In the case of practical reasoning whose conclusion is the formation of an intention, the agent is noninferentially disposed to respond reliably by producing suitable performances—which are qualified as actions by having such a provenance; these two sorts of evaluations appear in Davidson's account in the role of *premises* and of *conclusions*, respectively. The account offered here denies that what is expressed by the *prima facie* evaluative judgments Davidson understands as codifying pro-attitudes need appear as explicit premises in such reasoning; they correspond to the endorsement of a pattern of practical inference as entitlement-preserving. Such practical inferential commitments may be made explicit in the form of doxastic commitments with assertible contents (and so be available for duty as explicit premises) if suitable explicating vocabulary is available—just as theoretical inferential commitments may, but in general need not, be codified explicitly by the use of conditionals. In this use, then, normative expressions exemplified by 'should' as it appears in rules of conduct play an inference-explicating role on the practical side that is analogous to that played by 'if . . . then . . .' on the doxastic side;<sup>60</sup> in neither case does the omission of a premise that codifies a material propriety of inference, whether practical or doxastic, result in an enthymeme. Besides this *permissive* use of normative locutions such as 'should', which corresponds to Davidson's *prima facie* evaluative judgments, there is also a *committive* use, which corresponds to the "all-out" evaluative judgments that serve for him as intentions.

Recall the discussion of intentional explanation in Chapter 1. One of the ideas advanced there in connection with the suggestion that intentional states be understood in terms of deontic statuses and (propositional) attitudes toward them is that the conclusions of intentional explanations in the strict sense are *normative*, rather than *descriptive* claims. One attributes beliefs and desires (or other evaluations or pro-attitudes) and concludes from those attributions, not that the agent *will* perform an action of a certain kind, but that the agent is *committed* by those beliefs and desires to do so, that in the light of those other attitudes the agent *ought* (rationally) to do so. Intentional explanation illuminates what was done by showing why the intentional agent was *committed* to acting in that way. Under various circumstances it is possible to continue the inquiry and to ask why the agent acted in accord with that commitment. The response to such a question is not an intentional explanation, however, but a different sort of account—one showing why it is useful to offer intentional explanations of this individual, why treating the individual as a rational agent is a useful predictive and explanatory strategy. Explanations of this supplementary sort may appeal to how the organism is wired up and how it was trained so as to be able to respond reliably to the acknowledgment of a practical commitment by producing a performance of the sort specified in the content of that commitment. Such considerations are offstage from the point of view of intentional explanations proper, for these go only as far as showing what an agent should (rationally) do, what the agent is committed to do by the doxastic and inferential commitments that agent acknowledges.

Intentional explanations display sample pieces of practical reasoning, attributing theoretical and practical deontic statuses as premises and attributing a practical commitment as a conclusion. To serve as an intentional explanation of something the agent did, or to draw a conclusion about what the agent should do on which a prediction might be based about what the agent will do, these must be treated as commitment-preserving inferences. For the conclusion is that the intentional agent was or is *committed* to act in a certain manner. When the deontic scorekeeper attributes various commitments and concludes that therefore the agent *should* perform an action satisfying a particular description, the evaluative judgment expressed is of Davidson's second, all-things-considered kind.

So on the deontic scorekeeping approach there are two sorts of 'should', corresponding to the two sorts of evaluative judgment that Davidson considers: one involving *prima facie* evaluations suitable to serve as premises in practical reasoning, and one involving all-in evaluations suitable to serve as conclusions in practical reasoning. The first sort of 'should' is used to make explicit the endorsement, undertaken or attributed, of a pattern of practical reasoning, as in "Bank employees should wear neckties." The second sort of 'should' is used to make explicit commitments to act, which are attributed as the conclusions of committive practical inferences attributed in the course

of intentional explanation. To say this is to say that in the latter sort of use, 'should' expresses in the third person what is expressed in the first person by 'shall'. (Indeed, 'shall' and 'should' are etymologically linked in just the way suggested by this doctrine.)

Translated into the official scorekeeping terminology of deontic attitudes, this is the claim that while 'shall' is used to make explicit the *acknowledgment* (and therefore the *undertaking*) of a practical commitment to make some claim true, 'should' is used to make explicit the *attribution* of such practical commitments. The same piece of practical reasoning can be presented from either social perspective. In first-person, deliberative terms, the agent may acquire a practical commitment that would be made explicit (if the idiom encompasses sufficient expressive resources) by an overt utterance of "I shall wear a necktie," as the result of an inference from acknowledged commitments that would be made explicit (perhaps in response to a challenge to demonstrate entitlement to the practical conclusion) by an overt utterance of "I am a bank employee." In third-person, scorekeeping terms, the scorekeeper may attribute a practical commitment, adopting an attitude that would be made explicit (if the idiom encompasses sufficient expressive resources) by an overt utterance of "He should wear a necktie," as a result of an inference from attributed commitments that would be made explicit by an overt utterance of "He is a bank employee." The same piece of practical reasoning can be exhibited either by the one undertaking a practical commitment or by the one attributing it—significant either in deliberation regarding action or in assessment of such action. Indeed, deliberation—my considering various practical inferences in order to decide what I *shall* do—is just the internalization of assessment, the consideration of what anyone, given the relevant collateral commitments and circumstances (as they are taken by the assessor to be), *should* do.

### 6. *Weakness of the Will*

Of course 'should' has first-person uses as well. Some of these express only endorsement of patterns of permissive inference and so correspond to Davidson's merely *prima facie* evaluative judgments: "I should wear a necktie" (since I am a bank employee, but only if there is no better reason not to do so). But some also are self-attributions of commitments, in which one takes up a third-person perspective toward oneself, drawing conclusions about what one's reasons commit one to do: "I should (all things considered) drive to the airport." With practical commitments, as with doxastic ones, although acknowledging a commitment entails attributing it to oneself, the converse is not the case; though attributing a commitment to oneself is one way to *undertake* that commitment, this can be a consequential undertaking, rather than an acknowledgment. In particular, a self-attribution of the 'I should . . .' variety need not trigger the reliable differential dispositions to

respond to an acknowledgment of a practical commitment by producing a suitable performance. The noninferential significance of the deontic attitude that is made explicit by 'I should . . .' can be different from that of the deontic attitude that is made explicit by 'I shall . . .' Thus one can self-attribute a practical commitment without acknowledging it in the sense that matters for eliciting action.<sup>61</sup>

This possibility is one of the phenomena philosophers have discussed under the heading of weakness of the will, or *akrasia*—knowing the better and doing the worse. The attitude expressed by 'I should . . .' in its all-in sense does indeed commitment-entail that expressed by the corresponding 'I shall . . .' statement. But the difference between acknowledging a commitment and consequentially undertaking it depends on the fact that one does not always acknowledge the consequences of commitments that one acknowledges. In scorekeeping terms one can nonetheless be said to undertake those consequential commitments because the initial acknowledgment licenses others to attribute them. This distinction remains even when one takes up a third-person point of view toward oneself, as in deliberation about various possible courses of action when the agent traces out what commitments would be undertaken consequentially were certain others acknowledged. For in deliberating, an agent considers what commitments would be attributed by scorekeepers, under various circumstances. It is for this reason that one must be able to assess the conduct of others in order to deliberate about one's own.

The akratic's deliberations and intentions are out of step; the commitments acknowledged in the third-person theoretical way are incompatible with those acknowledged in the first-person practical way. The akratic agent is accordingly the analog on the practical side of the believer who undertakes incompatible doxastic commitments. It is one of the cardinal strengths of the deontic scorekeeping approach to intentional states in terms of normative statuses that there is nothing conceptually mysterious about the possibility of such incompatible commitments. Difficulties in coherently understanding akratic action and endorsement of incompatible beliefs arise from exclusive emphasis on a causal-functional model of intentional states.

The account of action presented here is a thoroughly Kantian one. For Kant, *will* is just the capacity for practical reasoning—that is, the capacity to derive performances from a conception of laws.<sup>62</sup> In the terminology introduced here, this is just the capacity to respond reliably to acknowledgments of commitments (the pragmatic version of "deriving from conceptions of laws") by producing suitable performances—suitable in terms of the way their descriptions line up with the contents of the practical commitments they either acknowledge or by the acknowledgment of which they are responsively elicited. For Kant the expressive role of 'ought' or 'should' (*Sollen*) is to make norms explicit in the form of imperatives. Specifically, such normative vocabulary "indicates the relation of an objective law of reason to a will

which is not in its subjective constitution necessarily determined by the law."<sup>63</sup> As construed here, normative vocabulary (of which 'ought' and 'should' are paradigmatic) has the logical expressive function of making explicit in the form of something that can be *said* (put in the form of a claim) an attitude that otherwise could be implicit only in what is *done*—namely, the endorsement of a pattern of practical reasoning. The propriety of a form of reasoning is the practical correlate of Kant's "objective (= valid, binding) law of reason," and its acknowledgment as constraining actual practical reasoning is its "subjective" relation to the will.

The rational will as described here is not a particularly puzzling phenomenon. Its normative dimension is explained by extending the account of discursive commitments to encompass not only doxastic but practical deontic statuses. Its causal dimension is explained by exploiting the analogy between discursive entries and exits, between action and perception. The relation between the normative and the causal aspects of rational willing or practical reasoning is explained by appealing to the causal efficacy of the deontic attitude of *acknowledging* commitments: acknowledgments of doxastic commitments can be reliably differentially elicited as responses to environmental situations in perception, and acknowledgments of practical commitments can reliably differentially elicit performances as responses in action. Reasons can be causes because deontic scorekeeping attitudes can play both normative and causal roles. There is much still to be learned about the empirical details of the differential responsive dispositions that make possible these discursive entry-and-exit practices, but it is not hard to understand in principle how there can be such things. We are rational creatures exactly insofar as our acknowledgment of discursive commitments makes a difference to what we go on to do—on the side of action, insofar as we incorporate a connection between what is expressed by 'should' and what is expressed by 'shall'.





---

**Abbreviations**

**Notes**

**Index**

---



---

# Abbreviations

---

The following abbreviations appear frequently throughout the Notes:

- BGS** Gottlob Frege. *Begriffsschrift*. 1879.
- "EPM"** Wilfrid S. Sellars. "Empiricism and the Philosophy of Mind." Reprinted in *Science, Perception, and Reality*. London: Routledge and Kegan Paul, 1963.
- FPL** Michael Dummett, *Frege's Philosophy of Language*. New York: Harper and Row, 1973.
- GL** Gottlob Frege. *Die Grundlagen der Arithmetik*. 1884. English translation, *Foundations of Arithmetic*, by J. L. Austin. Evanston, Ill.: Northwestern University Press, 1959.
- PI** Ludwig Wittgenstein. *Philosophical Investigations*. Translated by G. E. M. Anscombe. Oxford: Basil Blackwell, 1953.
- PPPW** J. Sicha, ed. *Pure Pragmatics and Possible Worlds: The Early Essays of Wilfrid Sellars*. Reseda, Calif.: Ridgeview Publishing, 1980.
- PW** H. Hermes, F. Kambartel, and F. Kaulbach, eds. *Gottlob Frege: Posthumous Writings*. Chicago: University of Chicago Press, 1979.
- RFM** Ludwig Wittgenstein. *Remarks on the Foundations of Mathematics*. Edited by G. H. von Wright, R. Rhees, and G. E. M. Anscombe. Translated by G. E. M. Anscombe. Cambridge: MIT Press, 1972.
- "STSSD"** Robert Brandom. "Singular Terms and Sentential Sign Designs." *Philosophical Topics* 15, no. 1 (Spring 1987): 125–167.
- "ÜSB"** Gottlob Frege. "Über Sinn und Bedeutung." In *Philosophical Writings of Gottlob Frege*, ed. P. Geach and M. Black, pp. 56–78. Oxford: Blackwell, 1969.



---

# Notes

---

## 1. Toward a Normative Pragmatics

1. The particular way in which Kant understands what theoretical and practical concepts are rules for doing—namely manipulating representations either by synthesizing many under one or as determinations of the will—depends on further, independent commitments that are not here in question.
2. That is the lesson of his “Was ist Aufklärung?” Indeed, for this reason the line between Cartesian and Kantian approaches should not be drawn so sharply as to imply that Descartes had no inkling of the significance of normativity, which becomes an explicit concern for Kant. His idea of the mental as a special stuff can be seen as a response to those issues, as yet only dimly appreciated. Descartes’s sense of the mental as special is precisely an inchoate awareness that its essence lies in rational, hence normative, interconnectedness. This makes it impossible to fit into what we now think of as nature, according to a conception of nature that was being formed around Descartes’s time. (Thanks are due to John McDowell for emphasizing this important point.)
3. In the unpublished 1897 draft of “Logic,” in *PW*, p. 147.
4. From another fragment on logic, *ibid.*, p. 4.
5. *Ibid.*, p. 145.
6. *Ibid.*, p. 144.
7. *Ibid.*, p. 145.
8. *Ibid.*, p. 128.
9. *Ibid.*, p. 4. Sometimes the point is put in terms of reasons, correct inference, or

justification: “Logic has a closer affinity with ethics [than psychology] . . . Here, too, we can talk of justification, and here, too, this is not simply a matter of relating what actually took place or of showing that things had to happen as they did and not in any other way” (ibid.).

10. Crispin Wright calls this Wittgenstein’s “contractual” model of meaning and understanding (though for reasons that will emerge, the overtones of explicitness incorporated in this way of talking about the normative dimension are less than happy) (*Wittgenstein on the Foundations of Mathematics*, [Cambridge: Harvard University Press, 1980], p. 19). John McDowell describes it as the idea that we are “committed to certain patterns of linguistic usage by the meanings we attach to expressions” (“Wittgenstein on Following a Rule,” *Synthese* 58 [1984] 325–363). As Saul Kripke puts it: “The relation of meaning and intention to future action is normative, not descriptive” (*Wittgenstein on Rules and Private Language*, [Cambridge: Harvard University Press, 1982], p. 37).
11. For example, at *PI*, § 146.
12. For example, at *RFM*, I121.
13. *PI*, § 193.
14. *RFM*, I116. One crucial difference is that the *laws* of society are explicit—they *say* what is correct and what is not. The assumption that all the laws of inference (not just those of logic) are explicit in this sense generates a regress, discussed in the next section.
15. *PI*, § 195.
16. *Ibid.*, § 217.
17. Two sorts of norms have been pointed out as involved in attributions of intentional states. On the one hand, intentional states stand in normative relations to each other: acquiring one belief commits one to believing its inferential consequences, intending to make-true a certain claim commits one to intending the necessary means, having certain constellations of beliefs and desires can commit one to form corresponding intentions, and so on. On the other hand, intentional states stand in normative relations to states of affairs that are not intentional states: there is a certain sort of normative *accord* between a belief and the state of affairs that must obtain for it to be true, between a desire and the states of affairs that would satisfy it, between an expectation and the states of affairs that would fulfill it, between an order and the performances that would count as obeying it. The first sort of normative relation is broadly *inferential*, the second is broadly *referential*. Although Wittgenstein invokes both sorts, his primary concern is with the latter. The strategy of this work is to start with the former kind of norm and to explain the latter kind in terms of it.
18. Thus the norms incorporated in the content of a belief concern not only what other beliefs one is committed to by having that belief (and in the context of other intentional states, how one is committed to act) but also how one thereby is committed to the world’s being—to be assessed by determining what objects one’s belief is *about*, and what is *true* of them.
19. The more general Kantian view at stake is that concerning the normative character of concept use. The more specific view is the understanding of norms as having the form of explicit rules. The juridical idiom he employs systematically obscures the distinction between these two commitments.
20. This is a different sense from the one that Sellars, whose views are discussed below, attaches to this expression.

21. *PI*, § 201.
22. *Ibid.*, § 84.
23. *Ibid.*, § 198.
24. *Ibid.*, § 201.
25. *Ibid.*, § 289 and *RFM*, V33.
26. *PI*, § 202.
27. *Ibid.*, § 199.
28. *Mind* 4 (1895): 278–280.
29. Gilbert Ryle, *The Concept of Mind* (New York: Barnes and Noble, Harper and Row, 1949), chap. 2.
30. *PI*, § 78.
31. What matters for the present project is the opposition between these two orders of explanation. But since they have been set out in connection with actual historical figures, as the lesson Wittgenstein has to teach Kant, it should at once be acknowledged, if only parenthetically, that when one looks at the details, Kant is somewhat better off than he appears in this sketch, for he does appreciate the point that Wittgenstein is making. Kant's acknowledgment of the possibility of a regress of rules appears in his discussion of the faculty of *judgment* (*Urteils-kraft*): "If understanding in general is to be viewed as the faculty of rules, judgment will be the faculty of subsuming under rules; that is, of distinguishing whether something does or does not stand under a given rule (*casus datae legis*). General logic contains and can contain no rules for judgment . . . If it sought to give general instructions how we are to subsume under these rules, that is, to distinguish whether something does or does not come under them, that could only be by means of another rule. This in turn, for the very reason that it is a rule, again demands guidance from judgment. And thus it appears that, though understanding is capable of being instructed, and of being equipped with rules, judgment is a peculiar talent which can be practised only, and cannot be taught" (*Critique of Pure Reason*, A132/B171). The regress-of-rules argument is here explicitly acknowledged, and the conclusion drawn that there must be some more practical capacity to distinguish correct from incorrect, at least in the case of applying rules. Very little is made of this point in the first two Critiques, however. Kant's own development of this appreciation of the fundamental character of this faculty of acknowledging norms implicit in the practice of applying explicit rules, in the third Critique, has an immense significance for Hegel's pragmatism, but only his formulation of the issue seems to have influenced Wittgenstein's. The Appendix to this chapter discusses Wittgenstein's use of 'rule' in more detail.
32. P. 60 of "Realism and the New Way of Words," in *PPPW*, pp. 219–256. Another early paper that is important in this connection is "A Semantical Solution to the Mind-Body Problem" (also in *PPPW*), which argues for the paired claims (1) that mental concepts are semantic, metalinguistic concepts and (2) that semantic concepts are normative concepts.
33. "Some Reflections on Language Games," in *Science, Perception, and Reality* (London: Routledge and Kegan Paul, 1963), p. 321.
34. *Ibid.*
35. From "Language, Rules, Behavior," *PPPW*, p. 155. In a similar vein he says: "The mode of existence of a rule is as a generalization written in flesh and blood, or nerve and sinew, rather than in pen and ink" (from the same essay, p. 139). Talk

of rules as generalizations, even incarnate ones, is dangerous in this connection, however, for it flirts with a reductive regularism (about which more below) that identifies proprieties of practice with regularities of conduct.

36. Sellars, "Some Reflections on Language Games," p. 322.
37. *Logic*, trans. R. S. Hartman and W. Schwarz (New York: Dover Publications, 1974), p. 3.
38. Kripke, *Wittgenstein on Rules*.
39. McDowell, "Wittgenstein on Following a Rule," p. 342. It should be acknowledged that McDowell construes the structure of Wittgenstein's argument differently from the way it is presented here. He takes it that the identification of understanding with interpreting presents two unacceptable alternatives: either the regress of rules does not stop, in which case the norms evanesce, since every action is in accord with any given norm on some interpretation and fails to accord on some other, or platonistic, self-applying norms are imagined as the *last* interpretation. Here this platonistic "rails laid out to infinity" misconstrual was presented as arising independently of identifying understanding with interpreting—as a way of misunderstanding norms on a quasi-causal model.
40. It should be clear that to insist on this point is *not* to claim that one *cannot* explicitly say what ought to be done, say by promulgating a rule or giving an order. Nor is it to claim that where one *does* follow such an explicit rule, one must be interpreting it. Precisely not. In the typical case the understanding of what is explicit, the following of a rule, is itself practical—the exercise of implicit understanding or "know-how." One of the central tasks of this work is to say what one must be able to *do* in order to count as in this sense understanding an explicit claim, rule, or order.
41. This usage of 'discursive' is Kant's. See for instance pp. 21, 34, 82 of his *Logic*.
42. Making out this distinction is really the subject of the whole of the *Critique of Practical Reason*. The rational will is defined this way in Section 7 of Part I, Book 1, Chapter 1, p. 32 of the Akademie Textausgabe.
43. "Heidegger on Being a Person," *Nous* 16 (1982): 16.
44. This point is related to McDowell's criticism, discussed below, of social regularity theories of the sort Kripke and Wright attribute to Wittgenstein, which make the community of assessors incorrigible.
45. Although this seems the natural way to elaborate the picture, it is not evidently incoherent to imagine one organism shaping its own behavior by responding to its responses with positively and negatively reinforcing behavior. What makes such a suggestion odd is that one would think that the capacity to distinguish correct from incorrect performance that is exercised in the postulated responsive disposition to assess would also be available at the time the original performance is produced, so that no behavior-shaping ground would be gained by the two-stage procedure. But this need not be the case; the assessment might be addressed toward the performance as characterized by its consequences, discernible more readily in the event than the advent. It is no doubt more difficult to tell a story about how such self-reinforcing patterns of behavior might come about in one animal than in a group, for the behavior-shaping in question is not here, as it is in the regulist versions, *deliberate*, a matter of explicitly expressible intentions. Yet the issue of what it would be for there to be norms implicit in practice ought to be kept distinct from the issue of how such practices might in



fact plausibly arise. If the intra-organism reinforcement story is coherent, then regularity versions of the sanctions approach to implicit norms need be social only in the sense that they essentially involve the distinction of perspective between producing performances and assessing them. This contrasts with Sellars's story, in which the behavior shaping by reinforcement is deliberate and the regularity of conduct aimed at is accordingly explicitly expressible by the assessors, even though the assessors and the assessed may be time-slices of the same organisms (and full-fledged membership in the community may require playing both roles at some time). That account seems genuinely to require that there be performances where the assessing individual and the individual producing the performance being assessed are distinct. For only cases of this sort can be appealed to in making intelligible the norms implicit in grasp of a concept in such a way as to have any leverage at all against the regress-of-interpretations argument that motivates this approach. So the diachronic regulist sanctions theory sees norms as implicit in specifically *social* practice in a stronger sense.

46. For instance, McDowell "Wittgenstein on Following a Rule," p. 350: "If regularities in the verbal behaviour of an isolated individual, described in norm-free terms, do not add up to meaning, it is quite obscure how it could somehow make all the difference if there are several individuals with matching regularities."
47. Construing communal assessment regularity theories (paradigmatically those Kripke and Wright attribute to Wittgenstein) as offering an account of what it is for norms to be implicit in practice is implicitly disagreeing with one of McDowell's central criticisms of the relevance of such theories to Wittgenstein's text. McDowell objects: "The fundamental trouble is that Kripke makes nothing of Wittgenstein's concern to reject the assimilation of understanding to interpretation" (*ibid.*, p. 343). He is right that neither Kripke nor Wright makes anything of this crucial motivating line of thought, which is rehearsed in Section II of this chapter. But he overlooks the fact that the theory they do elaborate can nonetheless be understood as an attempt to address just the considerations that are motivated by the regress-of-interpretations argument. For they can be seen as concerned to provide a notion of what it is for norms to be implicit in practice or for practice to be implicitly norm-governed rather than explicitly rule-governed. As will emerge, this repudiation is consistent with endorsement of McDowell's other criticisms of this line of thought and interpretation.
48. Kripke, *Wittgenstein on Rules*, p. 108.
49. Wright, *Wittgenstein on the Foundations of Mathematics*, p. 220. Although Wright explicitly addresses only the significance of linguistic performances, his point applies more generally to acting correctly according to one's intentional states. A communal assessment regularity theory is also put forward in the author's "Freedom and Constraint by Norms," *American Philosophical Quarterly*, April 1977, pp. 187–196, reprinted in *Hermeneutics and Praxis*, ed. R. Hollinger (Notre Dame, Ind.: University of Notre Dame Press 1985).
50. Davidson is a notable exception, taking linguistic practice and therefore intentionality to be essentially social only in the sense that it can be made intelligible only in the context of mutual interpretation—an *I-thou* relationship, in the current terminology.
51. One example of how this demand could be met by defining community membership in such a way as to preserve the distinction between those governed by

the practice and those whose practice it is they are governed by, without disjoining the groups, is provided by Sellars's account of pattern-governed practice. This is what going intergenerational does for him—the judgments of the assessors who train new community members are authoritative, and those they assess and train are the community members subject to their authority.

52. These correspond to the two sorts of objections to individual regularity or dispositional theories that Kripke (*Wittgenstein on Rules*) offers. McDowell ("Wittgenstein on Following a Rule") argues that the social regularity theory Kripke then suggests Wittgenstein endorses in response is subject to an objection of the second sort, namely that it fails to distinguish between a claim's being correct (normative status) and its being taken to be correct (normative attitude) by the community as a whole. It is argued above that this approach also falls foul by importing illicit notions of communal assessment, normative statuses such as community membership (being subject to communal authority), and expertise (exercising communal authority).
53. Some account along these lines has been a popular post-Enlightenment reading of what is being allegorically communicated by the supernatural retributive strand in Christian ethical theory.
54. Mill, in *Utilitarianism* (reprinted in *Essential Works of John Stuart Mill*, ed. Max Lerner [New York: Bantam Books, 1965], 3:215), introduces the vocabulary of internal and external sanctions, but to point to a different distinction than that intended here. His "internal" sanctions are internal to the individual (rather than to the space of norms). A paradigm would be feelings of shame or guilt.
55. *PI*, § 201.
56. Typically, though not in every case, by not letting it begin—since in the commonest cases we understand explicit claims, rules, principles, orders, and so on *without* interpreting them.
57. The theory developed in this work incorporates both of these suggestions. But at this point in the exposition no specific interpretation of either has yet been endorsed.
58. Wright (*Wittgenstein on the Foundations of Mathematics*) and Kripke (*Wittgenstein on Rules*) offer interpretations along these general lines.
59. This reading is closely related to McDowell's criticism of readings of passages such as these: "If I have exhausted the justifications I have reached bedrock, and my spade is turned. Then I am inclined to say: 'This is simply what I do'" (*PI*, § 217); "When I obey a rule, I do not choose. I obey the rule blindly" (§ 219); "How do I know [how I intend the pattern to be continued]?—If that means 'Have I reasons?' the answer is: my reasons will soon give out. And then I shall act, without reasons" (§ 211). These succumb to the temptation to conclude that, "at the level of 'bedrock' (where justifications have come to an end), there is nothing but verbal behavior" (McDowell, "Wittgenstein on Following a Rule," p. 341). That is to think of the bedrock of unreflective practice exclusively in nonnormative terms of behavioral dispositions and regularities. But as McDowell points out, one should not conclude that where justification has run out, normative assessments no longer apply. That is just what the regress-of-rules argument for the existence of norms implicit in practice shows. Wittgenstein says, in a claim important enough to appear verbatim in both *PI* and *RFM*: "To use the word without a justification does not mean to use it wrongfully [zu

*Unrecht gebrauchen*]" [RFM, V33; PI, § 289]. As McDowell says: "It seems clear that the point of this is precisely to prevent the leaching out of norms from our picture of 'bedrock'—from our picture, that is, of how things are at the deepest level at which we may sensibly contemplate the place of language in the world" ("Wittgenstein on Following a Rule," p. 341).

60. P. 5, par. 2, of *On the Law of Nature and of Nations*, trans. of 1688 ed. C. H. Oldfather and W. A. Oldfather, vol. 2 in the Classics of International Law series (reprint, New York: Oceana Publications; London: Wiley & Sons, 1964).
61. *Ibid.*, par. 3.
62. *Ibid.*, par. 4.
63. *Ibid.*, p. 6, par. 5.
64. One of Heidegger's central concerns in *Being and Time* (trans. J. Macquarrie and E. Robinson [New York: Harper and Row, 1963]) is to deny this characteristic Enlightenment thought, by describing how the value-free presence-at-hand (*Vorhandenheit*) studied by the physicist is abstracted from the value-laden readiness-to-hand (*Zuhandenheit*) of everyday life. Here is a characteristic statement: "In interpreting we do not so to speak, throw a 'signification' over some naked thing which is present-at-hand, we do not stick a value on it" (sec. 32, p. 190). This view is discussed in detail in the author's "Heidegger's Categories in *Being and Time*," *Monist* 66, no. 3 (July 1983): 387–409.
65. The evolution of physics from its "atoms in the void" conception has not appreciably altered the difficulty of fitting norms into the natural scientific world-picture. It is this difficulty that motivates both the Kantian dualism of norm and fact and the Kantian normative idealism that subordinates the latter to the former. Since the normative force of the better reason is not easily understood in terms of the sort of causal forces invoked by Newton, a normative conception of the way in which the necessity codified in laws outruns mere regularities is called in to support an understanding of causes in terms of proprieties governing the employment of concepts. Whatever one thinks of this heroic inversion strategy for reuniting the disparate elements of the Kantian dualism, its motivation underscores the difficulty of accommodating the normative within the natural.
66. It is somewhat disingenuous to characterize his view in terms of *our* attitudes. Although his primary concern in this work is with the moral attributes instituted by human beings, Pufendorf also acknowledges (as which seventeenth-century philosopher did not?) that God is also an intelligent being and can also impose or institute moral attributes. In this sense, God is treated as one of us. But even for God, in this respect *primus inter pares*, creation of the physical world is one thing, imposing moral attributes on it something else.
67. *Leviathan*, chap. 6, p. 24.
68. David Gauthier, *Morals by Agreement* (New York: Oxford University Press, 1986), p. 21.
69. He does not endorse a corresponding thesis for the merely prudential or instrumental norms according to a conception of which we are also capable of acting. But he also holds that there could not be a being that had a rational will in the sense of being able to act according to a conception of a prudential rule or maxim, but did not have a rational will in the sense of being subject to moral norms. So even though not all the rules we acknowledge or act according to

conceptions of are moral rules, we can still be demarcated as the ones who act according to moral rules, for which he does endorse a version of the thesis being discussed.

70. Kant acknowledges his most immediate debt to Rousseau. (It has seemed incongruous to some that a portrait of that wild, intemperate, irregular figure should have provided the sole adornment in the study of the excruciatingly continent and excessively rule-governed Kant.) This tradition is treated as the organizing theme of the Enlightenment in Kant's "Was ist Aufklärung?"
71. Pufendorf, *Law of Nature*, chap. 2, par. 6, p. 27.
72. *Ibid.*, chap. 5, par. 4, p. 89.
73. *Ibid.*, par. 9, p. 95.
74. *Ibid.*, par. 14, p. 107.
75. This is, of course, just as one would expect for an approach that takes its point of departure in construing norms from the example of explicit positive law. "So there are two parts of a law, one defining the offence, and one setting the penalty or the penal sanction; two parts, I say, and not two kinds of laws. For it is idle to say, 'Do this', if nothing follows; and it is equally absurd to say 'You will be punished', if the reason is not added, why punishment is deserved. It must, therefore, be borne in mind that the whole power of a law properly consists in its declaring what our superior wishes us to do or not to do, and what penalty awaits its violators" (*ibid.*). Thus the superior must have "the strength to threaten some evil against those who resist him" (par. 9, p. 95). Besides construing authority in terms of sanctions, Pufendorf also endorses two other central theses considered in the previous section of this chapter. For he takes it that by an obligation "we are bound by the necessity of doing something, for by it some moral bridle, as it were, is slipped over our liberty of action, so that we cannot rightly turn to any other quarter than that to which it directs. An obligation, however, can in no way so bind the will that it cannot, indeed, go contrary to it, although at its own peril" (chap. 5, par. 5, p. 90). Thus his conception of norms treats as essential the possibility of a distinction between what is in fact done and what ought to be done. Perhaps more remarkably, he develops his retributive picture of the practical expression of assessments by endorsing the idea of normatively *internal* sanctions: "an obligation affects the will morally, and fills its very being with such a particular sense, that it is forced of itself to weigh its own actions, and to judge itself *worthy* of some censure, unless it conforms to a prescribed rule . . . Again, an obligation differs in a special way from coercion, in that, while both ultimately point out some object of terror, the latter only shakes the will with an external force, and impels it to choose some undesired object only by the sense of an impending evil; while an obligation in addition forces a man to acknowledge of himself that the evil, which has been pointed out to the person who deviates from an announced rule, falls upon him *justly*" (*ibid.*, emphasis added).
76. *Ibid.*, par. 12, p. 101.
77. A contemporary version of this view—in particular of the sort of positive freedom (freedom to do new sorts of things, rather than freedom from constraint) that results from constraining oneself by specifically *linguistic* norms—is presented in the author's "Freedom and Constraint."
78. *PI*, § 258.

79. *Wittgenstein on the Foundations of Mathematics*.
80. *Ibid.*, p. 220.
81. McDowell, "Wittgenstein on Following a Rule," pp. 333–334.
82. Most of the discussion of Dennett refers to views propounded already in his "Intentional Systems," *Journal of Philosophy* 68, no. 4 (1971): 87–106; reprinted in *Brainstorms: Philosophical Essays on Mind and Psychology* (Montgomery, Vt.: Bradford Books, 1978); page references are to the reprint edition.
83. *Ibid.*, p. 221.
84. Davidson calls constellations of beliefs and pro-attitudes of this sort "primary reasons" for action, originally in "Actions, Reasons, and Causes," *Journal of Philosophy* 60 (1963), reprinted in *Actions and Events*, (New York: Oxford University Press, 1980), p. 4. This sort of intentional explanation is discussed in more detail in Chapter 4.
85. "Intentional Systems," p. 13.
86. *Ibid.*, p. 17.
87. John Searle, *Intentionality* (Cambridge: Cambridge University Press, 1983).
88. Berkeley: California University Press, 1969.
89. To say that the norms implicit in practices confer conceptual content is to say that having such content just consists in being governed by those proprieties.
90. *PI*, § 54. This catalog might be taken to refer only to what he calls "definite" rules, as also in *ibid.*, § 81.
91. As he does at *ibid.*, §§ 224 and 225.
92. For example, at *ibid.*, § 142.
93. At *ibid.*, § 198, but his better wisdom may be expressed rather in the converse proposition at § 85.
94. *Ibid.*, § 653.
95. For example, at *ibid.*, § 237.
96. *Ibid.*, § 199.

## 2. Toward an Inferential Semantics

1. This choice of terminology follows Hegel's use of *anerkennen* in his *Phenomenology*.
2. Franz Brentano, *Psychology from the Empirical Standpoint*, trans. D. B. Terrell, quoted on pp. 119–120 in *Introduction to the Philosophy of Mind: Readings from Descartes to Strawson*, ed. H. Morick (Glenview, Ill.: Scott, Foresman, 1970).
3. John Searle, *Intentionality* (Cambridge: Cambridge University Press, 1983), p. 2.
4. Robert Stalnaker, *Inquiry* (Cambridge: MIT Press, 1984), p. 2.
5. Searle, *Intentionality*, p. 17.
6. It is because of the distinction indicated by Brentano's reservation that the hyphenated phrase "object-representing" is used here, rather than more committal talk of "representing objects," to mark the categorial contrast with propositional contentfulness.
7. Cited by Roderick Chisholm, on p. 140 in *Introduction to the Philosophy of Mind*, ed. H. Morick (Glenview, Ill.: Scott Foresman, 1970).
8. "The regular connexion [*Verknüpfung*] between a sign, its sense, and its reference is of such a kind that to the sign there corresponds [*entsprechen*] a definite

sense and to that in turn a definite reference" ("ÜSB," p. 58). As an abbreviation only, Frege also allows talk of the expression, rather than the sense it expresses, designating or referring to what it represents: "To make short and exact expressions possible, let the following phraseology be established: A proper name [word, sign combination, expression] expresses its sense, stands for or designates its reference. By means of [mit] a sign we express its sense and designate its reference" (p. 61).

9. "Artificial Intelligence as Philosophy and as Psychology," in *Brainstorms: Philosophical Essays on Mind and Psychology* (Montgomery, Vt.: Bradford Books, 1978), p. 122.
10. It may be helpful in clearing up an incipient misunderstanding to remark here that in the official idiom to be developed and employed in this work, linguistic *expressions*, in the sense of marks and noises, do not need to be separately mentioned at this point. For it is not tokens but *tokenings* that are in the first instance considered as contentful. Sign-designs, the linguistic vehicles of content, are meaningful only at one remove, in virtue of their involvement in linguistic performances that express intentional states and attitudes.
 

The token/tokening distinction can often be overlooked (so that the theoretical decision as to explanatory priority alluded to here does not even arise) in the case of evanescent tokens such as utterances; the uttering/uttered ambiguity need not be resolved. The issue becomes more evident if one thinks about more permanent tokens, as when the religious enthusiast walks around the city with a sign in the shape of an arrow, inscribed "YOU are a sinner!" and points it at various passersby. In such a case the different tokenings have different contents, even though only one token is involved. The payoff of the policy mentioned here accordingly comes when the use of indexicals and other tokenings that are in principle unrepeatable becomes a topic, in Chapters 7 and 8.
11. *Critique of Pure Reason*, A97.
12. Bertrand Russell, in *The Philosophy of Logical Atomism*, ed. D. Pears (Lasalle, Ill.: Open Court, 1985), is a case in point.
13. *Critique of Pure Reason*, A69/B94.
14. *Ibid.*
15. *Ibid.*, A68/B93.
16. *Ibid.*, A69/B94.
17. *Ibid.*, A126.
18. *Ibid.*, A79/B104–105. The "transcendental element" introduced in this way is just reference to objects.
19. "Notes for L. Darmstädter," in *PW*, p. 253.
20. "Boole's Logical Calculus and the *Begriffsschrift*," in *PW*, pp. 16–17.
21. The concept of substitution and its significance in such a *decompositional* semantic program are investigated in detail in Chapter 6.
22. *GL*, sec. 60. The claim that "only in the context of a proposition [*Satz*] does a name have any meaning" is enunciated also in the Introduction (p. x), and in secs. 46 and 62.
23. *PW*, p. 232.
24. "ÜSB," p. 63.
25. *Ibid.*, pp. 57 and 58.
26. *Ibid.*, pp. 62–63.

27. *PW*, p. 144.
28. "My Basic Logical Insights," *PW*, p. 252. See also the opening pages of the "Logic" of 1897, beginning at *PW*, p. 129.
29. This is what Kant is getting at in seeing the "transcendental element" of referring to objects as introduced into representations by their role in judgment, in the passage quoted above.
30. Desiring a mouse or desiring relief from hunger are best thought of as elliptically specified desires that one have or eat a mouse, that one's hunger be relieved. This should become apparent in the initial discussion of practical reasoning in Chapter 4. In any case, it is sufficient for the point being made here that the contents of the corresponding beliefs must be specified by sentential clauses.
31. A clue that is exploited in the account of this relation endorsed further along can be gleaned from looking at how to make explicit what a theorist becomes committed to in taking one complex object (for instance, a map) to be a representation of another (for instance, terrain). The theorist is claiming that from a certain kind of fact about the representing object (corresponding to a privileged vocabulary for describing it), it is possible to *infer* a certain kind of fact about the represented object. Thus from the fact that the blue squiggly line passes between a round dot and a square one, it is possible to infer that there is a river between a city whose population is less than 100,000 and one whose population is greater than 100,000. (This is not to say that when the representational relation is acknowledged only implicitly in the practice of someone using a complex object as a representation of another, the practitioner must be able to state explicitly the premises and conclusions of these inferences. See below at Chapter 8, Section II, Subsection 4 [such cross-references are abbreviated hereafter as 8.2.4].)
32. It is worth pointing out that this is not a difficulty that automatically confronts any theory that invokes a special ontological category of propositions in its account of claiming, judging, and believing. Obviously such accounts can accommodate the special status of propositional contents. A theory such as Stalnaker's, which understands propositions as sets of possible worlds and construes the attribution of propositionally contentful intentional states in terms of the use of the structure of possible worlds to *measure* those states for the purpose of explaining actions, is not vulnerable to the charge of semantic nominalism, of being in thrall to the model of designation. Such theories need involve no inappropriate assimilation of propositionally contentful states, attitudes, and performances to representings thought of as naming what they represent. They can respect the primacy of the propositional. They can do so precisely because they begin with the idea of an utterance expressing a proposition or a state exhibiting a propositional content. The question that then arises is what *expressing* a proposition has to do with *representing* anything. An account is required in any case of the relation between propositional contentfulness and object-representing contentfulness (purporting to represent objects). But only obfuscation results from talking in addition of sentences not only as expressing propositions and beliefs as having propositional contents but also of their *representing* propositions.
33. Although the point is put here in terms of cognition, a parallel point can be made on the side of rational action. For Kant understands the rational will as a faculty

- of causally determining particular acts through the conception of a general rule (*Critique of Practical Reason*, sec. 7).
34. It is abstracted by a comparative analysis, the forerunner of Frege's substitutional or functional method of analysis of the conceptual contents of judgments, which is the concern of Chapter 6 of this work.
  35. *Phenomenology*, par. 109. The erotic theory of classificatory consciousness arises in the order of exposition of the *Phenomenology* as the introduction to the theory of self-consciousness. Heidegger's successor concept of understanding in terms of taking something as something in practice is discussed in the author's "Heidegger's Categories in Being and Time," *Monist* 66, no. 3 (July 1983): 387–409.
  36. To make this point is not to claim that Hegel's erotic model does not have more resources (for instance for funding a distinction between correct and incorrect taking of something to be food) than are made available in the inorganic case. Consideration of inanimate objects suffices for the contrast of interest here, however.
  37. P. 262 of "Inference and Meaning," reprinted in *PPPW*.
  38. This brief sketch can no more than gesture at the rich development of these ideas in the *Phenomenology*. A fuller discussion of this important chapter in the tradition inherited by the approach pursued here lies outside the scope of this work. (It will be pursued on another occasion.) The few cryptic characterizations offered here are intended to serve only as placeholders, whose significance will become somewhat clearer as the way in which material contents can be construed in terms of inference and incompatibility and expressed by means of logical vocabulary are filled in as this chapter and the rest of the work proceed.
  39. In the autobiographical sketch in *Action, Knowledge, and Reality*, ed. H.-N. Castañeda (Indianapolis: Bobbs-Merrill, 1975), p. 285.
  40. A detailed accounts of their efforts, understood along these lines, is offered in the author's "Leibniz and Degrees of Perception," *Journal of the History of Philosophy* 19, no. 4 (October 1981): 447–479; and "Adequacy and the Individuation of Ideas in Spinoza's Ethics," *Journal of the History of Philosophy* 14 (April 1976): 147–162.
  41. *BGS*, sec. 3. Frege's word *richtig* here is usually misleadingly translated as 'valid'. The discussion below of the relation between materially and formally good inferences is intended to explain why 'correct' is a better translation here.
  42. *Ibid.*, sec. 2.
  43. "Boole's Logical Calculus and the *Begriffsschrift*," *PW*, pp. 16–17.
  44. *BGS*, sec. 3.
  45. *Logical Syntax of Language* (London: Routledge, Kegan, Paul, 1964), p. 175. Sellars's reference is in "Inference and Meaning," *PPPW*, p. 266.
  46. *FPL*, p. 432.
  47. *Ibid.*, p. 433.
  48. As will become clear, the idiom of material *inference* is not to be understood in relation to the use of the so-called material *conditional*.
  49. "Inference and Meaning," in *PPPW*, p. 261.
  50. Reprinted in *Brainstorms* (Montgomery, Vt.: Bradford Books, 1978), pp. 10–11.
  51. "Intentional Systems," p. 11.
  52. *Ibid.*



53. "Inference and Meaning," in *PPPW*, p. 265.
54. *Ibid.*, p. 284. This talk about the "framework" of logical transformation rules is just one expression of the attitude toward the relation between formal and material inference considered here. It would not be underwritten by the approach endorsed below, where logical vocabulary is picked out by its expressive role and then used to derive a notion of formal validity from material correctnesses of inference.
55. *Ibid.*, pp. 270–271.
56. *Ibid.*, p. 273.
57. *Ibid.*, p. 274.
58. It should be noticed that the point being made here has nothing to do with the relation in mathematical logic between proof theory and model theory. In particular, it is *not* being claimed that one need be concerned only with the former, to the exclusion of the latter. The concepts of arithmetic cannot be fully specified by finitely stateable rules of inference. Nevertheless, we do grasp those concepts. But this is just to say that we do in fact understand their inferential significance. To make explicit the inferences that articulate the concepts of arithmetic, we must employ model-theoretic metalanguages. This fact in no way impugns the inferential conception of conceptual content; it merely shows that traditional proof-theoretic metalanguages are not sufficiently expressively powerful to make such inferential roles explicit. The additional (inferential) expressive power added by metalanguages that employ the traditional semantic vocabulary of truth, denotation, and satisfaction is discussed in Part 2 below.
59. "Language, Rules, and Behavior," in *PPPW*, p. 136 n. 2.
60. From "Concepts As Involving Laws, and Inconceivable without Them," in *PPPW*, p. 122. The remark of A. J. Ayer referred to is from p. 17 of *Language, Truth, and Logic* (New York: Dover Publications, 1952).
61. "Boole's Logical Calculus and the *Begriffsschrift*," in *PW*, pp. 12–13.
62. *Ibid.*, p. 13.
63. *Ibid.*, p. 46.
64. *BGS*, Preface, in *From Frege to Gödel*, ed. Jean van Heijenoort (Cambridge: Harvard University Press, 1967), p. 7.
65. *Ibid.*
66. "Boole's Logical Calculus and the *Begriffsschrift*," in *PW*, p. 16.
67. This is a reason to reject the quasi-Tractarian view according to which nothing can count as claiming or asserting (and so nothing can count as inferring) unless the repertoire already contains logical vocabulary, so that the simplest claiming (the making explicit of anything) already presupposes the whole of logic.
68. See n. 28 above.
69. It will emerge in Chapter 3 that *entitlement*-preserving inferences are also important. They correspond roughly to inductive inferences in the same way that commitment-preserving ones correspond to deductive inferences.
70. In his fragment "Logic," Frege seems to endorse this order of explanation. He says: "To make a judgment because we are cognisant of other truths as providing a justification for it is *inferring*. There are laws governing this kind of justification, and to set up these laws of correct [*richtigen*] inference is the goal of logic . . . It would not perhaps be beside the mark to say that the laws of logic are nothing other than an unfolding of the content of the word 'true'" (*PW*, p. 3).

71. Only the sentential logical connectives are being addressed here. Identity and quantification, which raise special formal and philosophical difficulties, are discussed in later chapters.
72. 'Extensional' in this context can be made sense of in purely substitutional terms, without having to appeal to the sort of representational concepts in terms of which it is usually explicated [see 6.2 below].
73. The original investigation is in the author's "Varieties of Understanding," in *Reason and Rationality in Natural Science*, ed. N. Rescher (Lanham, Md.: University Press of America, 1985), pp. 27–51. The treatment there is cleaned up, corrected, and substantially extended in Mark Lance's "Normative Inferential Vocabulary: The Explication of Social Linguistic Practice" (Ph.D. diss., University of Pittsburgh, 1988), where the relevant completeness results are proven. The most interesting logical systems result from a semantics that combines pragmatically conferred incompatibility relations with pragmatically conferred entailment relations.
74. The actual procedure defines the introduction of a connective only as the principal connective in a formula and defines how to eliminate only principal occurrences. Full generality is nonetheless assured by working recursively. It should be remarked that according to the approach developed here, the standard Gentzen-style definitions for logical connectives are still possible for conjunction and disjunction, but the expressive role of conditionals, negation, and many other bits of logical vocabulary requires that they be understood as having quite another sort of introduction rule.
75. *FPL*, p. 453.
76. Noninferential reports are discussed in more detail in Chapter 4.
77. The *empirical* contribution to conceptual content made by noninferential *circumstances* of application in *perception* and the *practical* contribution to conceptual content made by noninferential *consequences* of application in *action* are discussed in Chapter 4.
78. The significance of this sort of example is explored in the author's "Truth and Assertibility," *Journal of Philosophy* #73, no. 6 (March 1976): 137–149. Ingredient contents are discussed below at 6.1.2.
79. *FPL*, p. 455; the following passage is on pp. 453–454.
80. *Ibid.*, pp. 456–457.
81. *Ibid.*, p. 455.
82. *Ibid.*, p. 454. It should be noted that inferential conservativeness is a weaker condition than derivability of circumstances from consequences (or vice versa). Showing how to derive one aspect from the other, using logic or prior inferential commitments, is sufficient but not necessary for conservativeness. I am grateful to Michael Kremer for this point.
83. N. Belnap, "Tonk, Plonk, and Plink," *Analysis* 22 (1962): 130–134, commenting on A. N. Prior's "Runabout Inference Ticket," *Analysis* 21 (1960–1961): 38–39.
84. *FPL*, p. 454.
85. Jonathan Bennett suggested this illustrative anecdote.
86. *FPL*, p. 455n.
87. *Ibid.*, p. 358.
88. In Quine's *From a Logical Point of View* (Cambridge: Harvard University Press, 1953), pp. 20–46.

89. This discussion addresses only versions of the project of offering truth conditions that envisage employing *other* concepts than those expressed by the words appearing in the sentences for which one is offering truth conditions, so as to offer *substantive* explications of those concepts. If truth conditions are conceived *modestly*, so that one is allowed to specify the truth condition for the sentence "Luther was a Schwabian" as simply as that Luther was a Schwabian, then the consequences pointed to in the text do not arise.
90. It should be acknowledged that although the discussion of this chapter has been framed throughout in terms of a stark opposition between two complementary orders of explanation—the representationalist and the inferentialist—these alternatives are not exhaustive. Other possibilities include treating neither representation nor inference as explanatorily prior to the other. One might then go on to explain both in terms of some third notion, which is treated as more fundamental. Or one might eschew reductive explanations in semantics entirely and remain contented with describing the relations among a family of mutually presupposing concepts—a family that includes representation, inference, claiming, referring, and so on.
91. Recall from the discussion in 1.4 above that the most serious objection McDowell levies against the social-practice theories of norms put into Wittgenstein's mouth by Wright and Kripke is that they have no room for the idea of proprieties of concept use that the whole community could be wrong about. As Wright puts it, these theories jettison the intuitive "ratification independence" of concept use for the special case where the ratifying attitudes of taking particular candidate applications of concepts to be correct or incorrect are those of the community as a whole.
92. The obligations involved in this order of explanation have just been indicated. It was suggested above that the corresponding explanatory demands on the contrary directions of explanation pursued by the intellectualist about norms, the formalist about logic, and the representationalist about content are difficult to meet. The intellectualist about norms has trouble explaining the norms governing the use or application of rules, principles, claims, and concepts. The formalist about logic has trouble explaining nonlogical content. The representationalist has trouble explaining specifically propositional content and its grasp. Of course there are various strategies for meeting or evading these demands. The present assembling of reminders and considerations intends only to sketch an alternative; it does not pretend to offer all-purpose refutations of the various contrary explanatory strategies that might be adopted.
93. This is an important point for Kant as well. His terminology in the *Logic*, where one-premise inferences are called "immediate" and multipremise inferences are called "mediated," greatly influences Hegel's use of those central technical terms of the *Phenomenology*.

### 3. Linguistic Practice and Discursive Commitment

1. The same structure is exhibited even if sentences are not taken as the primitive interpreted expressions. For example, if the basic stipulated assignment is of objects to singular terms, and sets of objects to predicates (or sets of sequences of objects, for multiplace predicates), then the results of simple syntactic predi-

- cations may be assigned truth-values as derived interpretants accordingly as the objects (or sequences) corresponding to the term(s) appearing in the predication are or are not included in the sets corresponding to the predicates.
2. *FPL*, p. 413. The view Dummett expounds in this passage differs from the one to be developed in this work in that he is concerned only with the contentfulness of linguistic expressions, not with that of intentional states and attitudes more generally; he considers only truth-conditional semantic interpretation; and he does not make clear the essentially normative character of the linguistic practices that constitute the use or working of the language.
  3. A paradigm of the recognition of this promissory note implicit in, for example, possible-worlds semantics is David Lewis's "Languages and Language," in *Language, Mind, and Knowledge*, Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson (Minneapolis: University of Minnesota Press, 1975), pp. 3–35; reprinted in Lewis's *Philosophical Papers*, vol. 1 (New York: Oxford University Press, 1983), chap. 11. Lewis's account of what it is for semantic interpretants to be appropriately associated with expressions by the use of language turns on his notion of *convention*, which appeals to propositionally contentful *intentions* and beliefs and so is not suitable to be extended to an account of the pragmatics corresponding to the contentfulness of such intentional states. In *Inquiry* (Cambridge: MIT Press, 1984), Robert Stalnaker provides an account of what it is for sets of possible worlds to be associated as the propositional contents of intentional states such as belief, appealing only to the possibility of intentional interpretation of intelligent behavior.
  4. H. P. Grice, "Meaning," *Philosophical Review* 66 (1957): 377–388. Also see Stephen Schiffer, *Meaning* (New York: Oxford University Press, 1972).
  5. David Lewis, *Convention: A Philosophical Study* (Cambridge: Harvard University Press, 1969), and "Languages and Language"; Jonathan Bennett, *Linguistic Behavior* (Cambridge: Cambridge University Press, 1976); John Searle, *Speech Acts* (Cambridge: Cambridge University Press, 1969); and *Expression and Meaning* (Cambridge: Cambridge University Press, 1979).
  6. Davidson takes this model of communication to entail that sharing a language is a merely practical, hypothetical necessity—it is convenient for members of a linguistic community to use the same noises to express the same thoughts because it minimizes the need for explicit theorizing about the intentions with which the noises are produced. The expectations or customs of other speakers, however, have no authority over how any individual is correctly understood. What matters is how the speaker intends to be understood or interpreted. Davidson's subtle position is different from the others mentioned in this connection in important ways, however, as he does not take it that the contents of these communicative intentions can be made sense of antecedently, in abstraction from interlocutors' interpretation of one another. Put another way, the intention to be interpreted one way rather than another that Davidson rightly takes to be essential to the meaningfulness of ordinary discourse can be understood to be implicit rather than propositionally explicit, an intention-in-action rather than a separately individuable prior intention. That is, its involving such an intention can be conceived as an automatic compliment paid to a performance in virtue of the fact that were the issue to be raised and the speaker sincerely to *disavow* the intention to be understood in a certain way, that would be evidence

that it had been misinterpreted. Such a view is importantly different from the Gricean picture of meanings as imposed on utterances by the antecedently contentful intentions of speakers, although Davidson is not always careful to register the distinction.

7. Jay Rosenberg, *Linguistic Representation* (Dordrecht: D. Reidel, 1975), chap. 2.
8. This thought leads John Searle (*Intentionality* [Cambridge: Cambridge University Press, 1983]) to insist that the intentionality of intentional states, unlike that of linguistic expressions, must be *intrinsic*. That is, Searle endorses the view, characteristic of agent semantics, that “I impose Intentionality on my utterances by intentionally conferring on them certain conditions of satisfaction which are the conditions of satisfaction of certain psychological states.” (p. 28). Accordingly, he is committed to a strong distinction between the contentfulness of utterances and that of the intentional states they express: “Since sentences—the sounds that come out of one’s mouth or the marks one makes on paper—are, considered in one way, just objects in the world like any other objects, their capacity to represent is not intrinsic, but is derived from the Intentionality of the mind. The Intentionality of mental states, on the other hand, is not derived from some more prior forms of Intentionality but is intrinsic to the states themselves. An agent uses a sentence to make a statement or ask a question, but he does not in that way use his beliefs and desires, he simply has them” (pp. vii–viii). Thus, “That the belief has those conditions of satisfaction is not something imposed on the belief by its being used at all. A belief is intrinsically a representation in this sense: it simply consists in an Intentional content and a psychological mode . . . [It is false that] in order for there to be a representation there must be some agent who uses some entity as a representation. This is true of pictures and sentences, i.e. of derived Intentionality, but not of Intentional states” (p. 22). The doctrine that the intentionality or contentfulness of intentional states and attitudes is *intrinsic* is in some ways the correlate in the domain of philosophical semantics of the method of *stipulation* in formal semantics. The theorist takes it that intentional states are just the sort of thing that comes with an intentional (paradigmatically propositional or representational) content. This is Descartes’s strategy—the mental is distinguished as being naturally *about* other things (it is the sort of thing other things have *objective* reality in). The theorist can talk about the consequences of such contentfulness but cannot be expected to have anything substantive to say about what that contentfulness consists in, apart from those consequences. Thus Searle denies that it is possible to give an analysis of intentionality, taking it rather that “Intentionality is, so to speak, a ground floor property of the mind” (pp. 14–15).
9. Robert Stalnaker, *Inquiry* (Cambridge: MIT Press, 1984), p. 4.
10. *Ibid.*, p. 5.
11. *Ibid.*, p. 6.
12. In spite of the acknowledgment of these alternatives implicit in his disjunctive formulation, however, Stalnaker restricts his consideration (and his arguments against the linguistic approach) to resemblance theories. Thus he concludes his description of the linguistic approach: “It is not essential to the linguistic picture that every thinking creature be capable of outward speech or that every one of our thoughts be expressible in our public language. All that is essential

is that thought be explained by analogy with speech" (*ibid.*, p. 5). This is true of resemblance theories—but not relation theories—of the significance of linguistic practice to intentional states. His further subdivision of the linguistic picture is phrased so as to accord with the restriction of consideration to analogical theories, in which thought is understood on the model of speech. "The development of the linguistic picture leads in two quite different directions which emphasize different analogies between speech and thought. One hypothesizes a language of thought, which may be different from any language used for communication; the other argues for the dependence of thought on the social activities of speech" (pp. 5–6). But in fact thought or the possession of contentful intentional states might be taken to depend on the social activities of speech *either* because the contents of intentional states must be modeled on and so understood in terms of the contents expressed by public speech acts in general *or* because it is essential to a state's having such content that the state can issue in a speech act by which it is publicly expressed—that is, either according to a linguistic theory of intentionality structured by resemblance or according to one structured by relation.

13. "Thought and Talk," in *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984), p. 156.
14. This and the subsequent passages are from *ibid.*, p. 170.
15. *FPL*, p. 362.
16. Hartry Field, "Mental Representation," *Erkenntnis* 13 (1978): 9–61. This strategy was put forward by Sellars already in 1953, in his "Semantical Solution to the Mind-Body Problem" (in *PPPW*, pp. 219–256; see especially secs. IV–VI), and remained at the center of his approach thereafter.
17. *Beurtheilbarer Inhalt*, introduced in the *BGS*, sec. 2.
18. This is not to say that in the full-fledged language game any particular move that is taken to be correct by the practitioners—even *all* the practitioners—thereby counts as correct. (That this promissory note is eventually redeemed is demonstrated by the objectivity proofs presented below in 8.6.5.) One of the primary tasks of this work is to begin to explain the way in which the linguistic community can institute incompatibilities relating subpractices, which then constrain its own assessments where those subpractices interact—that is, the way in which noninferential reporting practices and deliberate actions performed as a result of practical reasoning conspire to confer *objective* empirical content on the concepts they are inferentially linked to—the way in which what is said can come to answer for its correctness not to the ones using the language but to what they use it to talk *about*.
19. (Formal) inconsistency is to (material) incompatibility as formal logical validity of inference is to material correctness of inference. In each case the former should be defined in terms of the latter, for the reasons discussed above in 2.4.4. (The official interpretation of the <sup>s</sup>scare quotes<sup>s</sup> employed here is presented below in 8.4.5.)
20. In order to get a more realistic model, a shell might be added around this practice. A man, for example, who routinely fails to fulfill promises might be held responsible not only for failing to recognize himself as having undertaken a commitment (not recognizing his entitlement to entitle others to rely on him) but also for abusing the community practice of promising by his repeated attempts to claim or pretend to an authority that he has then withheld. Holding

such a man responsible in this way might consist in shunning him or in beating him with sticks, as in the original case. It is possible, in other words, to combine a general, externally defined sanction, with a specific, internally defined one. The former becomes a shell around the latter. The general sanction might be an instance of holding responsible for any abuse of linguistic authority, which might eventually lead to expulsion from the linguistic community (being treated like a parrot or a pariah). (This addition was suggested by Michael Kremer.)

21. The language of scorekeeping is suggested by David Lewis's "Scorekeeping in a Language Game," reprinted as Chapter 13 of his *Philosophical Papers*, pp. 233–249. More is made of this notion below, in Section IV.
22. The responsibility characteristic of action and the authority characteristic of perception are discussed in Chapter 4.
23. Mark Lance noticed that this definition permits one to consider *asymmetric* incompatibility relations as well as symmetric ones. He exploits this possibility formally in his development (in "Normative Inferential Vocabulary: The Explicitation of Social Linguistic Practice" [Ph.D. diss., University of Pittsburgh, 1988]) of the sort of incompatibility semantics originally suggested in the author's "Varieties of Understanding," in *Reason and Rationality in Natural Science*, ed. N. Rescher (Lanham, Md.: University Press of America, 1985), pp. 27–51. *Inconsistency* is the formal correlate of incompatibility. It is a logical notion, to be understood in terms of *negation*. But what makes a bit of vocabulary express negation is itself to be understood in terms of its relation to material incompatibility. The negation of a claim is defined as its minimal incompatible, the inferentially weakest claim that is entailed (in the commitment-preserving sense) by everything that is incompatible with the original claim.
24. As the examples discussed in the previous section indicate, it is not impossible to have authority without responsibility, or responsibility without authority. As students of organizational behavior will attest, however, it is a basic principle of social engineering that the stability and effectiveness of a practice are undercut if the authority accorded to some practitioners outruns their corresponding responsibilities, or vice versa. Linguistic practice as here construed is well designed in this respect.
25. Kurt Baier, "Responsibility and Freedom," in *Ethics and Society*, ed. R. T. DeGeorge (New York: Anchor Books, 1966), pp. 49–84.
26. John Searle, *Speech Acts*, p. 96.
27. Justificatory practices depend on entitlement-preserving inferences. But commitment-preserving inferences are also entitlement-preserving (though not conversely). If anyone who is committed to *p* is thereby committed to *q*, the only case in which entitlement to *p* plausibly would not carry with it entitlement to *q* is one in which the interlocutor is precluded from entitlement to *q* by concomitant commitment to something incompatible with it. But if *p* commitment-entails *q*, anything incompatible with *q* is incompatible with *p*, so under the circumstances described, the interlocutor could not be entitled to *p*.
28. The justificatory mode of entitlement inheritance requires that one invoke claims with *different* contents, for otherwise the 'stuttering' inference, from *p* to *p*, would count as a justification of *p*.
29. This is, as will become clear in terms of the model, compatible with the possibility in the fully developed practice of an interpreter correctly taking the entire community to be wrong about what commitments they are entitled to—but

such a judgment will always be that they are in some sense wrong by their own lights, that is, wrong given how they have committed themselves to its being proper to settle such questions and assess the answers.

30. Lewis, "Scorekeeping in a Language Game," p. 236.
31. *Ibid.*
32. *Ibid.*, p. 237.
33. *Ibid.*, p. 238.
34. Chapter 9 (Sections II and III) discusses the essential role played in such objective constraints by the fact that the interpreter who attributes discursive deontic scorekeeping practices to a community can use nonscorekeeping vocabulary with an antecedent use in specifying those practices.
35. Lewis, "Scorekeeping in a Language Game," p. 239.
36. *FPL*, p. 361.
37. *Ibid.*
38. "To begin with" because of the contribution made to semantic contents by the role sentences play in noninferentially elicited (but inferentially articulated) observation reports and in their role in giving rise to actions.
39. It is assumed throughout (though this requirement could be relaxed) that incompatibility is a symmetric relation. Also, if it is assumed that commitment-preserving inferences are entitlement preserving, in the absence of incompatible defeasors it follows that if everything incompatible with  $q$  is incompatible with  $p$ , then the inference from  $p$  to  $q$  is good both committively and permissively.
40. The correctness of such an inference according to  $A$  depends not only on the commitments and entitlements to commitments that  $A$  attributes to  $B$  but also on the commitments that  $A$  undertakes. When the expressive resources are available for explicit challenges to reliability inferences, these background premises of  $A$  will be cited in justification of those inferences and their conclusions. Under those circumstances they will be cited as constituting *standard conditions* for  $A$ 's observational authority with respect to this sort of content. This point is adverted to in the text further along.
41. Such a mapping directly characterizes what corresponds to the appropriate consequences of application of the expression. The circumstances of appropriate assertion (according to the scorekeeper in question) can be recovered from the full mapping, however, as the subset of initial scores in which the consequences of assertion include attribution, not only of commitment, but of entitlement to the assertion (according to the scorekeeper in question). The way in which entitlement is attributed to noninferential reports (thereby treating the reporter as reliable) shows that this class of appropriate circumstances of assertion is wider than what would result from taking the assertion to be in order only when the prior score already included an attribution of entitlement.
42. Philip Kremer and Mark Lance present some fascinating results concerning the explicit codification of commitment consequences in the form of logical conditionals in "The Logical Structure of Linguistic Commitment, I: Four Systems of Non-Relevant Commitment Entailment" and "The Logical Structure of Linguistic Commitment, II: Relevant Commitment Entailment," both forthcoming in the *Journal of Philosophical Logic*.
43. If one is entitled to  $p$  and  $p$  commitment-entails  $q$ , one is entitled to  $q$ —any entitlement-defeating incompatibilities to  $q$  equally defeat entitlement to  $p$ .



44. One can be (taken to be) entitled to claims one is not (taken to be) committed to—these are conclusions one is entitled to draw but has not yet committed oneself to. In this way one may be entitled to each of two mutually incompatible claims, so long as neither has been endorsed and commitment to it undertaken. So one might have good inductive reasons for believing that the barn is on fire (smoke, the particular noises that would usually accompany it, and so on) and a different set of good inductive reasons for believing it is not on fire (the alarm has not rung, it is pouring rain, the barn was just inspected, and so on). Either conclusion by itself could be defended, though one would cease to be entitled to it if already committed to the conclusion of the other argument. Attribution of entitlement to the consequences of conjoining incompatible contents to which one is severally entitled (because not committed to either) is avoided by closing entitlements without commitment only under committive inferences (if not defeated by incompatibilities) and closing only entitlements to commitments actually undertaken (according to the one keeping score) under permissive inferences.
45. Under suitable conditions. Considerations having to do with the possibility of coercion, insincerity, shyness, and so on are systematically suppressed in presenting the model of assertional practice, on the grounds that they are intelligible only against a background of propositional contents conferred by the sorts of interactions considered here. Real-world phenomena such as these (which presumably will be present at every stage in the development of actual practice) create play between the proprieties of practice (especially scorekeeping practice) an interpreter takes a community to be bound by and their actual behavior. The status of this discrepancy—a normative generalization of the competence-performance distinction—is discussed in Chapter 9.
46. Daniel Dennett, "Intentional Systems," in *Brainstorms* (Montgomery, Vt.: Bradford Books, 1978), p. 19.
47. *Ibid.*, p. 20.
48. *Ibid.*
49. There are subtleties that require qualifying this formula, some of which are discussed below in 8.5.2.

#### 4. Perception and Action

1. The further conditions that have been suggested in response to examples of the sort Gettier first presented are not discussed here. Those impressed by the significance of these counterexamples to the sufficiency of JTB analyses of "S knows that *p*" may want to treat the term 'knowledge' as it appears in the rational reconstruction presented here as really invoking knowledge\* (which is defined simply as justified true belief). The considerations motivating many of the proposals for a fourth condition on knowing can be straightforwardly transposed into the idiom of this work, so those who believe that the status reconstructed here must be further specified in order to deserve the central role it is given might try the experiment of seeing how their favorite candidate looks in deontic scorekeeping guise. The basic elements of the social practice model's construal of the statuses corresponding to justification, truth, and belief do not turn on the kind of niceties concerning their interaction that attempts to for-

mulate a fourth condition must address. It is explained below why knowledge\* deserves to be accorded a fundamental explanatory role, regardless of its relation to what is expressed by the English word 'know'.

2. Of course, there is no necessity to adopt either direction of explanation; it may be that neither term is intelligible apart from its relation to the other.
3. This claim about the conceptually basic level of practices is compatible with the institution of a distinction in more sophisticated practices between *claiming* that *p* and *claiming to know* that *p*. This point is discussed further along.
4. The distinction being made here between the present approach and standard ones is different from, though intimately related to, the fundamental difference between understanding belief, justification, truth, and so knowledge as kinds of *normative status* rather than as kinds of *natural state*, so that one looks for *proprieties*, rather than *properties*, corresponding to them. The connection between the issues consists in the difficulty of appreciating the significance of the social distinction of attitude between acknowledging and attributing unless one is already thinking of what these are attitudes toward in normative terms, as commitments and entitlements—that is, deontic statuses, not descriptive states.
5. Being a *logical* being—having the expressive resources to make propositionally explicit crucial semantic and pragmatic features of the social practices in virtue of which one is a linguistic, rational, cognitive being—is a further, optional stage of development, which presupposes this fundamental one.
6. Insofar as it addresses sapience rather than sentience—that is, insofar as it is concerned with intentionality in the sense defined by possession of *propositional* attitudes.
7. One source of these views is the “thermometer view of knowledge” put forward by D. Armstrong in *Belief, Truth, and Knowledge* (Cambridge: Cambridge University Press, 1973). F. Dretske defends such a view in *Knowledge and the Flow of Information* (Cambridge: MIT Press, 1981), and more recently in “The Need to Know,” in *Knowledge and Skepticism*, ed. M. Clay and K. Lehrer (Boulder, Colo.: Westview Press, 1989). A. Goldman’s development of a reliabilist theory is expounded in *Epistemology and Cognition* (Cambridge: Harvard University Press, 1986), and again in “Précis and Update of *Epistemology and Cognition*,” in *Knowledge and Skepticism*, ed. Clay and Lehrer. R. Nozick puts forward a version of reliabilism in *Philosophical Explanations* (Cambridge: Harvard University Press, Belknap Press, 1981). M. Swain has a version that takes reasons more seriously than the others, in *Reasons and Knowledge* (Ithaca: Cornell University Press, 1981), and more recently in “Justification, Reasons, and Reliability,” *Synthese* 64, no. 1 (1985): 69–92. This list is intended only to provide a reasonably representative sample.
8. This sort of talk involves a promissory note—only in Chapter 6 is the official account developed to the point that it explains (in terms of substitution inferences) the relation between applying a *predicate* and making a *claim*.
9. Alvin Goldman, “Discrimination and Perceptual Knowledge,” *Journal of Philosophy* 73, no. 20 (1976).
10. J. L. Austin “Other Minds,” in his *Philosophical Papers* (New York: Oxford University Press, 1961).
11. Not objective “to begin with” because, as was acknowledged already in Chapter 1, it is a critical criterion of adequacy of any account of the use of empirical

concepts that it be able to explain how *in the end* objective proprieties governing that use can come into play—how claims can be understood as true or false regardless of whether anyone or everyone takes them to be so, depending rather on how things are with what the claims are about. Not until Chapter 8 will assembly be complete of the raw materials necessary to explain what is involved in this sense of objectivity. One important clue has been put on the table already, however. For it has been pointed out how a distinction can arise between what someone to whom a commitment is attributed is *justified* or *entitled* to believe, and what is in fact *true*. This is just the social-perspectival distinction of attitude between *attributing* deontic statuses and *undertaking* them. It is in terms of this fundamental *social* articulation of deontic attitudes that the possession by claims and concepts of *objective* representational content is eventually to be understood.

12. "EPM," sec. 32.
13. *Ibid.*, sec. 34.
14. *Ibid.*, sec. 35.
15. *Ibid.*, sec. 32.
16. *Ibid.*, sec. 35.
17. *Ibid.*
18. *Ibid.*, sec. 36.
19. This possibility was floated in the author's original discussion in "Asserting."
20. It is a commonplace among teachers of mathematics that students often profess to be completely unable to deal with problems of a certain sort long after they are in fact able to solve them reliably.
21. Assuming that the attributor of knowledge both attributes to Monique commitment to the claim and undertakes such commitment (that is, endorses the claim and so takes it to be true).
22. Sellars, "Some Reflections on Language Games," in Sellars, *Science, Perception, and Reality* (New York: Routledge and Kegan Paul, 1963).
23. "EPM," sec. 32.
24. "There are two ways in which a sentence token can have credibility: (1) The authority may accrue to it, so to speak, from above, that is, as being a token of a sentence type all the tokens of which, in a certain use, have credibility, e.g. ' $2 + 2 = 4$ '. In that case, let us say that token credibility is inherited from type authority. (2) The credibility may accrue to it from the fact that it came to exist in a certain way in a certain set of circumstances, e.g. 'This is red'. Here token credibility is not derived from type credibility" (*ibid.*).
25. "Epistemology Naturalized," in *Ontological Relativity and Other Essays* (New York: Columbia University Press, 1969).
26. As they are called in the author's "Asserting," p. 643. Cf. Hegel's remark in a related context that "one barren assurance is of just as much worth as another," in the Introduction to the *Phenomenology* (par. 76).
27. Recall from the earlier discussion that Stalnaker introduces the linguistic approach in a way that leaves room for relational linguistic accounts such as Davidson's and the one pursued here, but he then discusses and argues against only the reductive versions.
28. Though just how this latter possibility should be understood is not officially addressed until Chapter 8.

29. As will become clear, it would be another sort of reductive mistake to *identify* rational agency with preference-maximizing.
30. David Lewis, *Convention: A Philosophical Study* (Cambridge: Harvard University Press, 1969), also "Languages and Language," in *Language, Mind, and Knowledge*, Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson (Minneapolis: University of Minnesota Press, 1975).
31. "Communication and Convention," in *Inquiries into Truth and Interpretation* (New York: Oxford University Press, 1984), p. 280.
32. *FPL*, pp. 298ff., 354–356.
33. "Some Reflections on Language Games."
34. Though as will emerge below, in some cases the nonlinguistic intentional performance just *is* (has the scorekeeping significance of) the acknowledgment of a practical commitment, rather than being a *response* to such an acknowledgment. The difference corresponds to that between intentions-in-action and prior intentions.
35. Section V below discusses some different patterns of availability of practical reasons across interlocutors.
36. Here 'theoretical' is opposed to 'practical', as pertaining to relations exclusively between doxastic discursive commitments. This use ought not to be confused with the sense of 'theoretical' that is opposed to 'observational', within the doxastic sphere. In the latter usage, following Sellars's practice (in "Empiricism and the Philosophy of Mind" and elsewhere), theoretical claims are distinguished as those one cannot become entitled to noninferentially, by the exercise of reliable differential responsive dispositions to acknowledge doxastic commitments. Theoretical vocabulary is then distinguished as that which appears only in claims that are theoretical in this sense. It is this usage that stands behind Sellars's claim that the distinction between the observable and the theoretical is not ontological but only methodological. Neptune was a theoretical entity so long as claims about it could be arrived at only inferentially, as based on its perturbation of the observable orbits of other planets. It became observable, however, once we built telescopes powerful enough to make it subject to non-inferential reporting. Something is theoretical or observable in this sense only relative to our practices; nothing is "intrinsically" theoretical.
37. Davidson wants to analyze intention in terms of reasons. A representative formulation is "Someone who acts with a certain intention acts for a reason" ["Intending," reprinted as Chapter 5 of *Actions and Events* [New York: Oxford University Press, 1980], p. 84). Irrational actions accordingly pose a problem for him (which he addresses in "How Is Weakness of the Will Possible?" reprinted as Chapter 2 of *Actions and Events*). According to the present view, he is mistaking a global condition on intention for a local condition, as a result of failing to distinguish commitment from entitlement.
38. This formulation is intended to encompass both committive and permissive inferences. The latter can be thought of as conveying *entitlement* to the commitments that are their premises to *entitlement* to the commitments that are their conclusions. In scorekeeping terms, for an attributor to endorse such a permissive practical inference is to take it that anyone who is committed and entitled to the premises is entitled, though not committed, to the conclusion. Nonetheless, the conclusion can be thought of as a commitment, because what one is entitled to is in the first instance a commitment.

39. This use of 'shall' is based on Sellars's technical usage of 'shall' as an expression of intention. The treatment of action, intention, and practical reasoning presented here owes a great deal to Sellars's seminal work on the topic, "Thought and Action," in *Freedom and Determinism*, ed. Keith Lehrer (New York: Random House, 1966).
40. The point just made about the inheritance by practical commitments of inferential relations from the corresponding doxastic commitments can be illustrated in these terms by noting that if "I will wear a necktie" commitment-entails "I will wear something around my neck," then "I shall wear a necktie" commitment-entails "I shall wear something around my neck."
41. Beginning with "Actions, Reasons, and Causes," reprinted as Chapter 1 in *Actions and Events*. Quotations are from p. 4.
42. The expressive role of such ascriptions should be understood by analogy to that of the explicit ascriptions of doxastic commitment discussed below in 8.1–4.
43. "Intending" p. 86.
44. Hempel discusses this feature of inductive arguments in detail in *Aspects of Scientific Explanation* (New York: Free Press, 1965), pp. 394–403. An example he offers (recast in the idiom of this work) is that one can have good evidence both for the inference codified by the conditional "If the barometer falls, it almost certainly will rain" and for the inference codified by "If the sky is red at night, it almost certainly will not rain." Since there are occasions on which one can be entitled to commitment to both of the antecedents, the incompatibility of the conclusions shows that these inferences cannot be commitment preserving. But they can each be entitlement preserving, even though in the situation where one is entitled to both antecedents, assertion of either can serve as a challenge to the conclusion of the other inference.
45. Literally at the end of this story, in Chapter 9, building on the social-perspectival account of objectivity developed in Chapter 8.
46. "Intending," p. 84.
47. *Ibid.*, p. 85.
48. G. E. M. Anscombe, *Intention* (Oxford: Blackwell, 1959), and Davidson, originally in "Actions, Reasons, and Causes."
49. Davidson, "Actions, Reasons, and Causes," p. 8.
50. *Ibid.*, p. xiii.
51. This and the next passage quoted are from John Searle, *Intentionality* (Cambridge: Cambridge University Press, 1983), pp. 84–85.
52. Sellars, "Thought and Action," p. 110.
53. *Ibid.*, p. 109.
54. This is the conclusion of Davidson's "Intending," summarized at pp. 100–101.
55. Sellars's central text on this topic is "Thought and Action." The view is introduced in his earlier "Imperatives, Intention, and the Logic of 'Ought'," *Methodos* 8 (1956): 228–268, and is developed at greater length in the Tsanoff lectures, delivered at Rice University in 1978, entitled "On Reasoning about Values." Castañeda's treatment dates to his 1952 University of Minnesota master's thesis, "An Essay on the Logic of Commands and Norms." The fullest statement of his view is in *Thinking and Doing: The Philosophical Foundations of Institutions* (Dordrecht: D. Reidel, 1975). A concise summary can be found in "The Two-Fold Structure and the Unity of Practical Thinking," in *Action Theory*, ed. M. Brand and D. Walton (Dordrecht: D. Reidel, 1976), pp. 105–130.

56. Indeed, in his extremely useful summary of Sellars's views on this topic (pp. 149–188 in *The Synoptic Vision: Essays on the Philosophy of Wilfrid Sellars*, by C. F. Delaney, Michael J. Loux, Gary Gutting, and W. David Solomon [Notre Dame, Ind.: University of Notre Dame Press, 1977]), W. David Solomon paraphrases Sellars's view about the intentions expressed by his regimented 'shall' locutions in just this way: "Whereas expressions of intention manifest my commitment to act at some future time (perhaps precisely datable, perhaps not), volitions are commitments on my part to act here and now" (p. 163). Sellars's treatment of pro-attitudes is quite different from that presented here, and in any case, he has no general account of beliefs as commitments, to which intentions might be assimilated.
57. Davidson, "Intending," p. 90.
58. *Ibid.*
59. *Ibid.*
60. The analogy is not exact, for 'shall' indicates a kind of pragmatic force, the acknowledging of a practical commitment. A 'shall' statement is not an ordinary assertion in which that force becomes explicit as part of the content, for instance in the way in which 'believe' can be used to make the attribution of a doxastic commitment explicit as part of the assertible content of an ascription (as discussed in Chapter 8). As Sellars says: "'Shall', in spite of its logical role, can be said to be a *manner* rather than a *content* of thought" ("Thought and Action," p. 109). For 'shall' statements do not embed in more complex statements—paradigmatically as the antecedents of conditionals—in the way ordinary assertible contents must. Such embedding strips off the pragmatic force associated with the utterance of the embedded sentence. Thus when I say "If I believe that Kant liked turnips, then I believe that he liked some tuber," I am not saying that I believe either claim. The force of the (self-)ascription "I believe that Kant liked turnips" has been lost, but 'believe' still means exactly what it does in unembedded contexts, even though the pragmatic significance of uttering it is different. It is different with 'shall'. "I *shall* marry, so I won't be a bachelor" is a good inference, but when codified as a conditional, it takes the form "If I *should* marry, I won't be a bachelor," not the ungrammatical "If I *shall* marry, I won't be a bachelor," which contains a defective use of 'shall'. The rest of this section indicates why this shift from 'shall' to 'should' is required in embedded contexts. The result is that 'shall' statements do not make acknowledgments of practical commitments explicit as assertible contents in quite the same sense that 'believe' statements make attributions of doxastic commitment explicit as assertible contents.
61. Self-attribution of a deontic state is what is made explicit by certain uses of 'I' and is not equivalent to merely attributing it to oneself. For the latter may be attributing it to someone who is, as a matter of fact, though one is not aware of this fact, oneself—as I might attribute a certain commitment to whoever wrote the words appearing on a scrap of paper I find in the street, not realizing that I wrote them a year ago. It is not at all surprising that one can attribute a practical commitment to oneself without acknowledging it; it is more surprising that one can even self-attribute such a commitment without acknowledging it (though not without undertaking it). This issue is discussed in more detail below in 8.5.2.

62. *Foundations of the Metaphysics of Morals*, trans. Lewis White Beck (Indianapolis: Bobbs-Merrill, Library of Liberal Arts, 1959), p. 29; p. 412 of vol. 4 of the Akademie Textausgabe.
63. *Ibid.*, p. 30; p. 413.

## 5. The Expressive Role of Traditional Semantic Vocabulary

1. In an extended sense: intentions have propositionally expressible conditions of satisfaction corresponding to the claims the agent is practically committed to making-true.
2. Permissive or entitlement-preserving inferences, which are of the first importance in justificatory practices—particularly inductive ones—are notoriously difficult to parse in terms of truth conditions. It is because of the hope that the notion of *reliability* can supply what is wanted that that concept assumes the significance it has for epistemologists who understand the contents of knowledge claims in terms of truth conditions. That notion, and the gerrymandering difficulties to which it is subject, are discussed in Chapter 4.
3. “The Concept of Truth in Formalized Languages,” in *Logic, Semantics, Metamathematics: Papers from 1923–1938* (New York: Oxford University Press, 1956), pp. 152–278.
4. Thomas Gray, “Elegy Written in a Country Churchyard.”
5. An analogy is the way in which, once the expressive role of ‘looks’ or ‘seems’ is properly understood, it becomes apparent that the incorrigibility of statements formed by the use of these operators is trivial in a way that makes them unsuitable precisely for the role of foundation of knowledge that Descartes assigned to them. The analysis required for this argument is presented in Sellars’s “EPM” and is sketched below in Section II.
6. Part of what is distinctive about the present approach, however, is that what are here treated as *semantic* primitives are themselves explained in terms of a prior *pragmatics*, which in turn appeals to *normative* primitives, themselves made available by mapping the theoretical idiom onto our ordinary talk.
7. Indeed, that is why semantic theorists, as opposed to linguists, have been in general so little interested in this notion—which will be taken here to be of absolutely the first importance to semantic theory. Chastain (whose work is cited below in Section IV) is a notable exception, as is Hintikka.
8. The material presented in this section and the next originally appeared as “Pragmatism, Phenomenalism, and Truth Talk,” in *Realism (Midwest Studies in Philosophy 12 [1988])*, pp. 75–93.
9. The point here does not concern merely the *senses* of the contrasted expressions, but the *extensions* they determine. The appropriateness of this question would have to be defended by adducing cases in which a belief apparently “worked” and was not true, or vice versa. Such cases are not far to seek. This sort of argument is considered more carefully below.
10. *Pragmatism* (Cambridge: Harvard University Press, 1978), bound with its sequel, *The Meaning of Truth*, which as here interpreted ought to be titled *The Meaning of Taking-True*. For an important assessment on a larger scale, see R. Rorty’s “Pragmatism, Davidson, and Truth,” in *Philosophy of Donald Davidson: A Perspective on Inquiries into Truth and Interpretation*, ed. E. LePore (Oxford:

---

# Index

---

- < > type designation, 314  
// token designation, 314  
[ ] corner quotes, 688n25, 704–705n24  
s s scare quotes, 545–547, 588–590, 672n19
- 'about', xvii, 138, 500–503, 520, 547–548, 590. *See also* aboutness  
aboutness, 6, 70, 279, 280, 306, 333, 336, 514, 518, 528, 547–548, 568, 583, 622, 647. *See also* representation  
and *de re* ascriptions, 391, 500, 502, 503, 547–548, 584  
and objectivity, 138, 530, 649, 672n18  
strong and weak, 584–586, 706n45  
abstraction, 419–422, 445, 614, 695n34  
abstract objects. *See* objects, abstract  
acknowledgment, 31–32, 35, 37, 43, 50–52, 202, 262–266, 525, 541, 542, 552–553, 629, 630, 633, 680n61. *See also* attributing; deontic attitudes; undertaking  
and attribution, xiv, xx, 55, 554, 646  
causal efficacy of, 259–262, 271, 596  
of commitments, 259, 262–266, 271, 596, 704n21  
explicit/implicit, 32, 624, 629, 639  
and undertaking, 194, 554  
action(s), xv, xx, 7, 8, 16, 69, 131, 142, 167, 230, 239, 243–245, 254–256, 290, 332, 507, 527–529, 614, 645. *See also* causes; deontic scorekeeping; making-true; normative vocabulary; practical commitments; practical reasoning; reliable differential responsive dispositions  
and cognition, 5, 200, 276, 289, 618, 665n33  
as language exit transition, 5, 221, 233–235, 294  
and linguistic practice, 155, 232  
and perception, 233–238, 261, 335, 336  
and practical reasoning, 231, 245–249, 255–256, 678n29  
and rationality, 230–233, 244–245  
and reasons, 7, 56, 171, 194, 203, 243–245, 249–256, 262  
and trying, 290–291, 295, 527–529  
adicity, 369, 385, 394, 408  
agentives, 556, 702n14  
agent semantics. *See* semantics, agent  
akrasia, 269–271. *See also* practical reasoning  
'A' language, 555–559, 708n65. *See also* 'I'  
anaphora, xvi, 285, 301–303, 306, 311–331, 438, 456, 462, 467, 471, 473, 486–494, 535–536, 559–567, 574, 580, 583, 585, 589, 622, 623, 683n35, 684n38, 698n80, 700n100, 705n30, 714n13.



anaphora (*continued*)

See also inference, substitution, and anaphora; substitution  
 ascription-structural, 516, 609–611, 701n7  
 in ascriptions, 536, 539, 577–578  
 as asymmetric token-recurrence structures, 455, 457, 467, 473, 490–491  
 and communication, 458, 459, 485–486, 503  
 and coreference, 306, 314, 485–486, 588  
 and deontic scorekeeping, 432, 460, 488, 494, 592, 633  
 expressive role of, 458, 467, 473, 588  
 inheritance of, 283, 303–304, 325, 391, 472–473  
 and inheritance of substitutional commitments, 198, 283, 325, 432, 455, 472, 473, 539  
 presupposed by deixis, 458, 462–468, 473, 511, 573, 585, 621  
 and quantification, 321, 491, 493  
 and recurrence, 309–310, 432, 453, 464, 490–492, 498, 621, 685n40  
 and reference, 306, 307, 322  
 and rigidity, 468–472, 698n80  
 anaphorically indirect definite descriptions, 305–307, 310–323, 464, 512, 547, 588, 589, 714n14  
 anaphoric antecedents, 301–305, 313–315, 450, 456, 459, 460, 486, 491–494  
 anaphoric chains, 307–309, 313–315, 319, 460, 462, 469, 470, 491, 493–494, 581–582  
 anaphoric commitments, 456, 472, 495, 583, 650, 706n33. *See also* substitution; token-recurrence structures  
 anaphoric dependents, 305–307, 309–331, 456, 458, 467, 470, 494, 574, 579–583  
 anaphoric initiators, 308–310, 458, 459, 462, 573–574  
 animals, 7, 33, 47, 83, 629, 691n42. *See also* intentional stance, simple  
 intentional states of, 150, 152, 155  
 Anscombe, G. E. M., 255–256, 558, 707n51, 708n69  
 Aoun, Joseph, 699n85  
 ascription(s), xx, 502, 506, 516, 530–531, 533–535, 537–539, 541–544, 547–548, 550, 566, 575–579, 581, 588–590, 595–596, 605, 609–613, 629, 703n19, 704n23, 711n102. *See also de dicto as-*

*criptions; de re ascriptions; explicit; pragmatics; representation*  
 as attribution and undertaking of commitments, 504–505, 531–533, 544, 640  
 direct discourse, 533–535, 542  
 of essentially indexical beliefs, 559–567  
 expressive role of, 503–504, 543, 588–589, 591, 608, 641, 679n42  
 indirect discourse, 534–535, 537, 539–540, 542  
 iterated, 545, 549, 608–613, 701n6, 702n13  
 paratactic theory of, 535–538, 577  
 of propositional attitudes, 279, 453, 485, 498–501, 530  
 regimentation of, 112, 501–503, 505–506, 516, 535, 540, 542–544, 546, 609, 613, 629  
 ascriptionally indirect definite descriptions, 547, 565, 590  
 assertibility, 121, 124, 299, 604  
 assertion(s), xiv, xxi, xxiii, 82, 156–159, 167–175, 179, 186, 190–191, 200–203, 278, 329, 335, 531, 639, 667n67, 675n45. *See also* deontic scorekeeping; doxastic commitments; giving and asking for reasons; linguistic practice; propositional contents; sentence(s); taking-true  
 authority/responsibility, 171–175, 179, 229, 234, 384, 532  
 bare, 228  
 and belief, 153–154, 194, 200, 229, 230  
 and commitment, 142, 179, 191, 200, 230  
 and communication, 156, 175–176, 221  
 and entitlement, 179, 180, 191, 200, 221, 229  
 and inference, 158, 167–168, 171, 339, 367, 687n11  
 and linguistic practice, 156, 158, 167–168, 173–176, 190, 200, 203, 221, 358, 623, 628–629, 637, 686n1  
 and pragmatic significance, 157, 168, 190, 200, 234, 358, 637  
 and truth, 199–204, 226, 229, 232  
 assertional commitments. *See* doxastic commitments  
 assessment, 18, 33, 37–42, 52, 64–66, 78, 236, 516, 526, 528, 614, 658n45, 659n47, 660n52. *See also* norms  
 appraisal/deliberation, 18–19, 287

- of correctness, 29, 32, 52–55, 63, 647  
and sanctions, 34–36, 42–45
- attitudes. *See* deontic attitudes; normative attitudes
- attributing, 196–197, 563, 598, 629, 630, 640, 642, 644, 645. *See also* deontic attitudes
- and acknowledging, xiv, xx, 55, 554, 646
- self-, 680n61
- and undertaking, 62, 161–163, 554, 596–598, 677n11
- Austin, J. L., 209, 288
- authority, 10–11, 50–52, 160, 163, 175, 212, 217, 226, 234, 238, 242, 677n24. *See also* inheritance of entitlements; justification; responsibility
- of noninferential reports, 214–221, 223–227, 234
- of norms, 50–52
- and responsibility, xii, 162–163, 165, 242, 673n24
- of testimony, 39–41, 175, 205, 234, 709n75
- autonomy, 20–21, 50–52, 171, 216, 265, 277, 293, 295, 465, 631
- auxiliary hypotheses, 92, 121–122, 475, 477, 488, 490, 517, 541, 596, 633, 635, 646. *See also* perspectives
- and communication, 139, 480–481, 511, 632
- and inferential significance, 139, 475, 478, 488, 504, 506, 509, 587, 633, 635
- and social perspectives, 519–520, 598, 608
- avowals, 194–196, 540, 577. *See also* assertion(s); belief(s)
- Bach-Peters sentences, 493–494
- Baier, Kurt, 172
- baptism, 580
- bare assertions. *See* assertion(s), bare
- barn facades, 209–212
- Bedeutung, 80–81, 111, 352
- behavior, 8. *See also* action(s); speech acts
- belief(s), xv, 5, 13–14, 17–18, 56, 142, 146, 152–157, 194–196, 201–202, 215, 228–230, 232–233, 250, 290–291, 515, 552–559, 573–579, 594–595, 656n18. *See also* assertions; doxastic commitments; intentionality; interpretation
- avowals of, 146, 195–196, 227–229, 540, 577
- de dicto/de re* [*see de dicto* ascriptions; *de re* ascriptions]
- essentially indexical, 551–552, 559–567
- mere, 229
- notional/relational, 547–548, 550
- object-dependent, 58, 551, 567–568, 570, 578, 583
- strong *de re*, 547–552, 555, 561–562, 566, 567, 581, 582, 585, 708n70, 709n88
- two senses of, 195, 507 (*see also* commitment(s), acknowledged/consequential)
- belief-forming mechanisms, 207–208, 210, 217
- 'believes', 116, 228, 261, 352, 498. *See also* belief(s)
- Belnap, Nuel D. Jr., 125, 300–301, 702n14
- Bennett, Jonathan F., 147, 668n85
- Boghossian, Paul, 326–327
- Boyd, Richard, 481
- Bradley, F. H., 375
- Brentano, Franz, 68, 70–71
- Burge, Tyler, 551, 566, 713n11
- Camp, Joseph, 300–301
- canonical designators, 112, 440, 442–449, 550, 569, 696n61, 706n41, 715n21. *See also* singular terms
- and existential commitments, 442–447
- Carnap, Rudolph, 96, 123, 215
- Carroll, Lewis, 22, 100–101, 206
- Castañeda, Hector-Neri, 263, 552, 563–564, 566
- categories, 335, 361, 362, 370–372, 404, 406
- causes, 11–14, 209–211, 259–262, 271, 427, 428, 429, 617, 620, 626, 687n2. *See also* action(s); norms; perception
- c-command, 700n106
- censure, 34, 37, 43
- certainty, 9–11
- challenges, 178, 192–193, 238
- Chastain, Charles, 307–309, 311, 456, 491, 683n32, 699n84
- Chomsky, Noam, 365–366
- circumstances and consequences of application, xiii, 18, 28, 40, 51, 89, 98, 117–131, 136, 159, 162, 182, 243, 331–332,

- circumstances and consequences of application (*continued*)  
 372, 383, 419, 421, 432, 433, 482, 541, 600, 618, 631, 640, 668n82, 674n41, 709n91, 714n12. *See also* concepts; inferentialism  
 and inference, 118–120, 126, 129, 131, 206, 225  
 'claims', 539–542, 711n96, 714n16. *See also* 'says'  
 claims, xv, 96, 112, 141, 153–157, 167, 229, 276, 291, 330, 453, 536, 586, 601, 619, 625, 641, 676n3. *See also* assertion(s)  
 and facts, 327–331, 622, 624, 625, 704n19  
 classification, 33, 35, 79, 85–86, 88–90, 666n35  
 and concepts, 85–89, 614–618, 623  
 codification, 107–108, 145, 267, 385, 397, 402, 403  
 cognition, 5, 200, 276, 289, 363, 614, 618, 665n33 *See also* deontic scorekeeping; doxastic commitments; knowledge; sapience  
 commands, 51, 64–66, 172, 270  
 commitment(s), 8, 55, 74, 157–161, 163–166, 168–172, 200, 201, 203, 209, 233, 250, 275, 474, 481, 583, 627, 649, 650, 674n43, 680n56. *See also* anaphoric commitments; deontic attitudes; deontic scorekeeping; deontic status(es); doxastic commitments; existential commitments; expressive commitments; inferential commitments; practical commitments; recurrence commitments; simple material substitution-inferential commitments (SMSICs); substitutional commitments  
 acknowledged/consequential, 193–197, 259–262, 270, 596, 633, 646, 704n21  
 attributing, 161–163, 205, 596–597, 646  
 and entitlement, xvi, 55, 159–161, 165, 245, 606, 675n44  
 making explicit, 334, 402, 505  
 undertaking, 161–163, 205, 596–597  
 communication, 473–475, 477–483, 485, 528, 566–567, 588, 635, 636, 644–647, 716n37. *See also* information; linguistic practice  
 and anaphora, 282, 458, 459, 477–482, 485–486, 503, 528  
 and assertion, 156, 174–176, 221, 357  
 and auxiliary hypotheses, 139, 480–481, 511  
 and contents, xxiii, 158, 169–170, 474, 478–481, 513–517, 562, 590, 633–636  
 and deontic scorekeeping, 156, 158, 169–170, 282, 474, 478, 480–481, 485, 496, 588, 635  
 and *de re* ascriptions, 508, 513–517, 519, 547, 701n7  
 and inheritance of entitlement, 175–176, 204, 205, 217, 226  
 and perspectival character of content, 485, 509–510, 635–636, 647  
 and strong *de re* beliefs, 561–562, 566  
 community, 594, 600, 630. *See also* *I-thou/I-we* sociality; *we*  
 and explicit discursive scorekeeping stance, 54, 639, 642–643  
 membership, 4, 39–40, 67, 643, 659n51, 660n52  
 and norms, 37–43, 67, 599–600, 639, 649  
 components, 298–299, 338, 343, 367, 376  
 compositionality/decompositionality, 351, 354–358, 361, 367  
 computers, 7  
 concepts, 8, 10–11, 18, 46, 79, 85–86, 89–91, 109, 113, 125–130, 207–208, 293, 484, 614, 616–620, 622–624, 634–636, 694n31. *See also* circumstances and consequences of application; content(s); grasping; inferentialism  
 conceptual change, 125, 127, 375, 402  
 conceptual schemes, 616, 618, 645  
 criticism of, 124, 126  
 dualistic conceptions of, 614–618  
 empirical, 119–120, 432, 458, 473  
 form/matter, 616–619, 622  
 inferential conception of, 87–91, 331, 551, 614–623, 633–635  
 and objectivity, 37, 109, 124, 126, 592–597, 614, 616, 624, 633, 636  
 practical, 9–10, 293, 616–618, 620, 622, 623  
 sharing, 587, 617, 620, 623, 631–636  
 theoretical, 9–10, 91, 616–617, 622  
 conceptual content(s), 4, 13, 107–108, 133, 225, 282, 426, 475, 582–583, 591, 636. *See also* conferring content; content(s); inferentialism  
 and discursive practice, 13–14, 107–108,

- 453, 485, 562, 590, 591, 614, 636, 645, 686n55
- in Frege, 95–96, 281, 348, 355
- inferential conception of, 95–96, 98, 102–106, 115, 130–132, 478, 485
- objectivity of, 63, 497, 529, 552–554, 593–594, 601–607, 630, 636
- perspectival character of, 485, 586–592
- and representation, 54, 140, 280, 583, 601–607, 630
- conditionals, 22, 103–104, 109, 112, 121–122, 247–248, 267, 352, 382, 393, 395, 396, 404, 414, 474, 498, 587, 666n48, 691n42, 697n70. *See also* explicit; logical vocabulary
- and embedding, 298–299
- expressive role of, xix, 108–110, 126–127, 383, 385, 395, 401–403, 498
- and inference, 98, 100, 101, 351, 352, 380, 381, 474, 602–603, 640, 693n56
- two-valued, 111–116, 344, 353, 530, 666n48
- conferring content, xv, 61, 64, 77, 91, 107–108, 115–117, 133, 137, 140, 145–147, 149, 151, 156, 261, 263, 592, 642, 663n89. *See also* content(s); practice(s)
- conceptual, 167, 169, 625, 627, 630, 632, 646, 649
- propositional, 7, 61, 63, 77, 141–142, 153, 159, 221, 275, 277, 284, 623, 638, 645
- conformism, 34
- conjunction, 115, 118, 352
- connectives, 117–130, 345, 347
- consciousness/self-consciousness, 85–87, 559, 614, 643, 644. *See also* classification; sapience; sentience
- consequences of application. *See* circumstances and consequences of application
- conservativeness, 123–125, 127, 129, 668n82
- constraints, 331–333, 529, 614, 617, 620, 621, 713n11, 714n11
- content(s), 67–71, 127, 137, 141–142, 145–147, 151, 153, 157, 167–175, 186–191, 199, 212, 330, 339, 343, 346–350, 352, 354, 359, 364, 368, 370–372, 374, 386, 401, 481, 516, 562, 605. *See also* conceptual content(s); conferring content; deontic scorekeeping; explicit; propositional content(s); representational contents; semantics
- and circumstances and consequences of application, 118, 129
- and discursive practice, 6, 12, 77–79, 199, 334, 359
- empirical (cognitive), 6, 212, 221, 225, 234, 295
- and force, 111, 186–190, 298–299
- freestanding, 338–340, 345, 348, 353, 354, 356, 358, 359
- inferential conception of, 89–90, 94, 130–132, 134–136, 186–190, 336, 342, 345, 347–348, 351, 353, 354, 374, 482, 485
- ingredient, 122, 338–340, 342, 345, 355, 356, 358, 359
- narrow/wide, 526, 703n18
- objective, 18, 136–139, 151, 153, 156, 595–596
- practical, xv, 221, 234, 295
- and representation, 70–72, 130–132, 135–138, 151, 153, 156, 517
- content-specifying expressions, 524, 646–647
- in ascriptions, 504–505, 588, 591, 608–611
- de dicto/de re*, 513–514, 516–517, 535, 543–544, 547–548, 584–585
- and social perspectives, 520, 595–596
- contexts, 316, 318, 342, 344, 347, 398–399, 483, 617. *See also* deixis; deontic scorekeeping
- extensional/intensional, 279, 281, 392
- heterogeneous/homogeneous, 344–346, 350, 356
- inferentially inverting, 381–382, 393, 394, 397, 403, 693n56
- sentential, 343–344, 346
- contract, 49–51, 242, 656n10
- conventions, 232, 233, 670n3
- coreference, 306, 308–309, 314, 318, 468, 485–486
- corporeality, 332, 631–632. *See also* solidity of practices
- correctness, 17–18, 32, 207–208, 278, 280, 291, 594–595, 614, 627–628, 632, 637, 666n41, 672n18. *See also* norms; pragmatics, normative; normative vocabulary
- assessment of, 9–11, 13–14, 29, 32, 52–55, 63, 647

- correctness (*continued*)  
 and practice, 22, 25, 32, 625, 628  
 and regularities, 27–28, 62, 207–208, 212  
 and rules, 8, 18–19, 24, 64–66, 628  
 and taking as correct, 25, 32–34, 37, 52–53, 63, 291, 626
- correspondence, 291, 326, 330–333
- cotypicality, 221, 308–310, 319, 469, 533, 535–536, 579, 622
- counterfactuals, 484, 574, 634–635, 648, 697n70
- counting, 438, 444
- criteria of identity, 416, 437, 438
- Davidson, Donald, 97, 150–152, 231, 232, 246–247, 255–256, 426, 523, 524, 614, 616, 622, 710n95  
 on ascriptions, 535–539, 550, 566, 575, 577  
 on interpretation, 15, 155, 262–264, 412, 599, 629, 659n50, 670n6, 678n37, 699n86, 716n37  
 on practical reasoning, 230, 246–248, 253–254, 259, 269, 337, 663n84
- de dicto* ascriptions, 485, 490, 502–508, 526, 529–534, 539–542, 566, 589, 601, 646, 703n19, 709n88, 710n95. *See also* ascription(s); belief(s); commitment(s); *de re* ascriptions and intentional explanation, 522–527  
 and ‘that’, 506, 598, 608
- deductive inferences. *See* inferences, commitment-preserving
- default-and-challenge structure of entitlement. *See* entitlement, default-and-challenge structure of
- defeasors, 191, 602–605
- deference, 204, 234, 453. *See also* authority, of testimony; communication; inheritance of entitlements, interpersonal, intracotent
- deferrals, 192–193, 196, 212, 531–532, 534, 536, 704n23
- definite descriptions, 298, 316, 386, 387, 415, 420, 431, 433, 439, 441, 456, 460, 464, 468, 469, 486, 617, 714n15. *See also* anaphorically indirect definite descriptions; ascriptionally indirect definite descriptions; singular terms and anaphora, 305, 308–309, 314, 458, 459, 464, 472  
 attributive/referential use of, 488, 579, 581  
 existence and uniqueness conditions on, 415, 434–435, 471  
 and indefinite descriptions, 307–310, 459
- definitization transformations, 309, 456, 459, 491–494
- deflationism, 325–327, 329
- deixis, 198, 282, 432, 453, 464, 468, 620.  
*See also* noninferential reports and anaphora, 306, 456, 458, 462, 464–468, 473, 511, 573, 585, 621
- deliberation, 158, 287, 290
- demarcation, 3–11, 46, 50–51, 87, 114, 135, 200, 644, 645. *See also* rationality; sapience; we
- demonstratives, 132, 282, 319, 439, 459–464, 466, 510–513, 550, 563, 582, 617, 621, 698n79, 705n30  
 and anaphora, 304, 462, 467, 469, 472, 511, 621
- Dennett, Daniel C., 15, 55–62, 73, 99–100, 195, 230, 547–548, 550
- denotation, 317, 318, 364, 547–548, 566, 706nn39,45. *See also* reference; representation
- ‘denotes’, 318. *See also* denotation
- deontic attitudes, 137, 221, 271, 290, 339, 497, 595–597, 599–607, 612–613, 637, 645, 649, 677n11. *See also* acknowledgment; attributing; deontic scorekeeping; deontic status(es); intentional states; normative attitudes; undertaking  
 and deontic statuses, 165–166, 357, 598–601, 623, 648, 687n11  
 hybrid, 187, 202, 220, 228, 297, 521–522, 525, 528–529, 545, 556, 568–569, 586, 704n20
- deontic scorekeeping, xiv, xvii, 141–143, 166–167, 181–187, 192–193, 202–203, 233–234, 242, 260, 264–266, 278, 282, 325, 330, 349, 436, 443, 461, 470, 480–481, 495–496, 503, 516, 555–560, 574, 584–585, 591, 605, 608, 624, 636, 639–642, 645, 648–649, 696n46, 700n99, 714n12. *See also* deontic attitudes; deontic status(es); entitlement; *I-thou/I-we* sociality; incompatibility; linguistic practice; pragmatics  
 and anaphora, 432, 460, 470, 487–488, 494

- and assertion, 157–159, 167–175, 190–191, 200–203
- and commitments, 157–166, 178–180, 193–196, 237, 243, 263
- and communication, 156, 174–176, 282, 478, 633
- and content, 141–142, 145–147, 167–175, 186–191
- and entitlement, 159–166, 176–180
- and interpretation, 475, 508, 630, 644–645
- model, 166, 168–172, 180–186, 190–191
- and objectivity, 324, 529, 601–607, 627–628, 636
- as perspectival, 185, 332, 488, 590, 602, 604, 627, 649
- and practical reasoning, 230–233, 244, 256, 625, 640
- and pragmatic significance, 142, 167–168, 182–190, 262, 284, 710n91
- and speech acts, 142, 182–193
- deontic status(es), 55, 165–166, 189–190, 201–205, 226, 237, 275, 290, 595–596, 636–637, 649. *See also* commitment(s); deontic attitudes; deontic scorekeeping; entitlement; intentional states; normative status(es); norms
- and deontic attitudes, 161–162, 165–166, 334, 357, 593, 598–601, 623, 648, 687n11
- and objectivity, 197, 201, 599–607
- de re* ascriptions, 391, 499–508, 511–515, 525, 542–548, 565, 566, 584, 589, 598, 601, 646, 648, 701n7, 703n19, 709n88, 710n95. *See also* ascription(s); *de dicto* ascriptions; information; representation
- denotational sense of, 547–548, 566, 706nn39,45
- expressive role of, 138, 502, 512, 516, 522, 525
- and intentional explanation, 522–527
- and objectivity, 528, 595, 598, 600–601
- and 'of', 506
- of propositional attitudes, 187, 279–281, 485, 490
- and representation, xvii, 279, 499–503, 517, 519, 520, 586
- strong, 503, 530, 566, 569–570, 581
- weak, 503, 513, 529–530, 547–552, 566
- Descartes, René, 6, 9–11, 31, 74, 93–94, 279, 614, 623, 655n2, 671n8, 681n5
- designatedness, 340–350, 352, 353, 356–358, 687n8. *See also* multivalues; substitution
- designation, 69, 84, 85, 665n32, 696n61. *See also* canonical designators; proper names; rigidity; semantics; singular terms
- desires, 56, 58, 240, 246, 250, 256
- Devitt, Michael, 481
- Dewey, John, 289, 299
- direct discourse ascriptions. *See* ascription(s), direct discourse
- direct reference. *See* reference, direct
- disavowals, 192–193, 670n6
- discursive attitudes. *See* deontic attitudes
- discursive commitments. *See* commitment(s)
- discursive practice. *See* linguistic practice
- discursive scorekeeping. *See* deontic scorekeeping
- discursive statuses. *See* deontic status(es)
- disjunction, 115, 429–430, 434–435, 438, 441–442
- dispositions, 28–29, 35, 42, 45–46, 208, 625, 628, 629, 636, 638. *See also* causes; regularities; reliable differential responsive dispositions
- disquotation, 300–301, 303, 323, 575, 577, 579, 581
- distinction
- acknowledging/attributing, 55, 193–197, 259–262, 270, 554, 596, 633, 646, 649
- action/perception, 7–8, 119–120, 209–211, 233–238, 261, 335, 336
- analytic/synthetic, 145, 345, 358, 484
- anaphoric/causal-historical approaches to proper names, 308–309, 458, 459, 470–471, 572–574, 579–583, 585–586
- asserting/infering, 158, 167–168, 171, 339, 347–348, 350, 351, 353, 355, 358, 359, 367
- asymmetric/symmetric, 372, 376–381, 384–385, 388, 391–395, 403, 455, 457, 490, 499, 564, 619–622
- attitude/status, 33, 37, 161–162, 165–166, 194, 197, 261, 290, 334, 357, 497, 593, 595–596, 598–601, 623, 648
- attributing/undertaking, 61, 62, 161–166, 196, 506–507, 525, 554, 596–601, 608
- authority/responsibility, xii, 161–165, 171, 174, 179, 229, 238, 242, 532

distinction (*continued*)

- Cartesian certainty/Kantian necessity, 9–11, 30, 636
- causal-functional/normative models of intentionality, 9, 15, 22, 30, 56–57, 60, 234, 270
- causal/normative, 12–15, 27–30, 33, 45–50, 93, 160, 196, 234, 617, 621, 625, 626
- circumstances/consequences of application, xiii, 18, 28, 40, 51, 89, 98, 117–131, 136, 159, 162, 182, 243, 331–332, 372, 383, 419, 421, 432, 433, 482, 541, 600, 618, 631, 640
- claimed/claiming, 327–330, 333, 488, 595, 606, 622–625
- cognitive/practical, xv, 6, 9–10, 119–120, 212, 221, 225, 234, 293, 295, 432, 458, 473, 616–623
- commitment/entitlement, xiv, 55, 159–161, 165–168, 179, 237, 238, 245, 252, 259–260, 541, 606, 649
- complex/simple predicates, 371, 406–409, 434, 436
- conferral/stipulation, 115–117, 145–147
- conferring/instituting, 46–50, 52–55, 107–108, 115–117, 137, 145–147, 165–167, 626–628, 638, 646
- correct/taken-correct, 29, 32, 52–55, 63, 595, 597, 647
- de dicto/de re* ascriptions, 499–508, 511–515, 529–534, 539–548
- deference/inference, 175–177, 192–193, 196, 204, 212, 218, 234, 241–242, 453, 531–536
- derivative/original intentionality, 60, 143, 171, 629–644
- designatedness values/multivaluedness, 340–350, 356–358
- distinction/dualism, 614–624, 626
- doing/saying, 30, 62, 77, 108–110, 115, 135, 639–641
- doxastic/practical commitments, 171, 233, 236, 238–243, 271
- entries/exits, language, 142, 221, 233–235, 258, 271, 335–336, 528, 632
- epistemically strong/weak *de re* ascriptions, 503, 513, 529–530, 547–552, 566, 569–570, 581
- explicit/implicit, xviii, xx, 18–26, 77–78, 85–86, 107–110, 121–122, 126–128, 147, 149, 247–249, 262–264, 432–433, 451–452, 498–499, 503–506, 512–513, 518–519, 531–532, 601–607, 612–613, 641, 649–650
- facts/norms, 58, 137, 331, 623–626
- force/content, 111, 186–190, 298–300, 322, 327–328
- formal/material inferences, 97–102, 104–105, 117, 133–136, 345, 383, 616, 619, 622
- formal/philosophical semantics, 143–145, 199
- freestanding/ingredient content, 122, 338–340, 342, 345, 348, 353–359
- inferentialism/representationalism, 6, 31–32, 92–94, 135–136, 205, 283, 285, 334–338, 360
- intentions in action/prior intentions, 256–259, 558
- intracontent/intercontent, intrapersonal/interpersonal, 169–170, 175–176, 179, 226, 241–242
- I-thou/I-we* sociality, 39, 62, 508, 522, 526, 590, 593, 598–607
- knowing-how/knowing-that, 23, 25–26, 101, 110, 135–136, 591, 641
- linguistic/pragmatic approaches to intentionality, 16, 22, 76, 148–150, 152, 229, 230, 631
- 'looks'/'is' talk, 292–297, 681n5
- making-true/taking-true, 5, 8, 13, 46, 233, 236, 277, 287–291, 297–299, 521
- natural/normative, xiii, xv, 12, 31, 35, 63, 149, 208, 289, 299–300, 624
- objective/subjective, 52–55, 197–198, 212, 526, 592–602, 604, 609
- practices/rules, 20–26, 32, 45, 55, 62, 64–66, 91, 99–100, 110, 509, 625
- purported/successful representation, 6–7, 70–75, 89–90, 360
- regularism/regulism, 18–29, 32, 36–42, 46, 62–63, 110, 208, 594
- reported/reporting tokenings in ascriptions, 535, 537, 541, 566
- sapience/sentience, 4–8, 87, 88, 231, 275–277, 520, 559, 591, 644
- 'shall'/'should', 245, 258–259, 261, 263–264, 267–271, 553
- true/taking-true, 287–292, 296–299, 322
- weak/strong/hyper-inferentialism, 131–132

- doing, 4, 80, 87, 91. *See also* explicit; knowing-how/knowing-that; pragmatics  
and saying, 30, 62, 77, 108–110, 115, 135, 639–641, 658n40
- donkey sentences, 490–493
- Donnellan, Keith, 488
- doxastic commitments, 142, 157–159, 167, 178–180, 200, 228, 238, 266, 276, 344, 346–348, 351, 472, 520, 679n40, 687n11. *See also* ascription(s); assertion(s); deontic scorekeeping; entitlement
- and assertion, 142, 194, 230
- and belief, 157, 196, 201, 228
- and practical commitments, 171, 233, 236, 238–243, 271
- Dretske, Fred I., 428, 430
- 'dthat', 469, 573
- dualism, 614–624, 626
- Dummett, Michael, 16, 17, 98–99, 116–118, 127, 144, 159, 187, 232–233, 296, 329, 347–348, 350, 352, 359
- on assertion and belief, 153–155, 200
- on circumstances and consequences of application, 117–131, 136, 162, 182
- on complex/simple predicates, 406, 688n27, 689n30
- on content, 339–340, 343, 345, 349, 358
- on inference/truth, 96–97, 111
- Edelberg, Walter, 691n48
- embedding, 298–300, 322, 338, 359, 381, 605, 609, 680n60. *See also* content(s); force; logic, multivalued  
test, 298–300, 604
- empiricism, 10, 85–86, 89–90, 614. *See also* representationalism
- "Empiricism and the Philosophy of Mind," 89–90, 215, 293, 465
- endorsement, 208, 293–294. *See also* assertion(s); commitment(s), undertaking
- Enlightenment, 6, 10, 47–49, 92, 93, 660n53
- entailment. *See* incompatibility; inference(s), commitment-preserving; inference(s), entitlement-preserving
- enthymemes, 98, 101, 206, 218, 635. *See also* material inferences  
in practical reasoning, 246–248, 252
- entitlement, 159–161, 176–179, 206, 208, 212–215, 221, 226–227, 238, 239, 334, 681n2. *See also* commitment(s); deontic attitudes; deontic scorekeeping; incompatibility; inheritance of entitlements  
and commitment, xiv, 55, 142, 159–161, 238, 245, 606, 675n44
- default-and-challenge structure of, 176–179, 184, 204, 221, 226, 238–239, 242
- and justification, 174, 204–206, 218, 532–533
- and reliability, 167, 206–211
- entitlement-preserving inferences. *See* inference(s), entitlement-preserving
- entries and exits. *See* language(s), entry/exit transitions
- epistemology, 6, 93, 201, 203, 207–208, 216
- equivalence classes, 281, 342, 345, 348, 351, 375, 392, 394, 400, 405, 408, 420–422, 450, 453, 454, 486, 619
- error, 31–32, 240, 295, 603, 606. *See also* ignorance
- evaluative judgments, 267
- Evans, Gareth, 460, 491, 567, 582, 696n62, 708n72
- existential commitments, 71, 304, 415, 416, 436, 440–449, 569, 694n17, 697n67, 706n41
- and canonical designators, 320, 442–444, 448
- negative existential claims, 319–320, 448
- and substitutional commitments, 431, 434, 440, 441, 445, 447
- exit transitions. *See* language(s), entry/exit transitions
- experts, 39–41, 600, 660n52, 709n75
- explanatory strategy, 45, 55–56, 84, 136, 154, 199–201, 403, 716n37
- assertion and judgment, 151, 199, 200, 202, 221, 232
- believing and claiming, 153–156
- bottom-up, 337–338, 340–341, 357, 358, 364
- inferentialism/representationalism, xvi, 6, 69, 93, 94, 97, 135–136, 149, 334–338, 495–496, 500, 503, 519, 584, 667n70, 669nn90,92
- intention and convention, 232, 696n46
- linguistic practice and rational agency, 155, 232



- explanatory strategy (*continued*)  
   and material inference, 132–133  
   naturalistic and normative, 149  
   normative phenomenalist, 25, 625, 628, 636–637  
   pragmatist/regulist, xiii, 25–26, 101, 112–113, 135, 149, 205, 657n31  
   semantic, 149, 199, 362, 495–496  
   sentential/subsentential expressions, 82–83  
   top-down, 337–338, 354, 358, 364  
 explicating locutions. *See* logical vocabulary
- explicit (making explicit), xviii, xx, 18–26, 32, 46, 62, 77, 85–86, 101, 112, 126–128, 140, 261, 279, 330, 395, 415, 418, 432–433, 442, 498, 505, 520, 530, 563, 586, 592, 599, 604–605, 624, 629, 639, 640, 644, 649–650. *See also* assertion(s); expressive rationality
- ascriptions, 228, 282, 498–499, 503–506, 530, 543, 588, 593, 612–613, 616, 629, 640, 679n42, 680n61
- deferrals, 226, 531–532, 534, 704n23
- de re* ascriptions, 280, 391, 508, 512–513, 515, 522, 584, 586, 595
- inferential commitments, xix, 104, 106, 108–110, 231, 247–248, 396, 667n58
- logical vocabulary, 107–108, 116, 319, 374, 382, 383, 399, 402, 418, 419, 498–499, 601–607, 619, 639
- norms, 18–23, 130, 247–249, 270, 271, 625, 639, 714n21
- ourselves, 275, 587, 641, 642, 650 (*see also* we)
- practical commitments, 247–249, 259, 262–264, 266, 267, 269, 403, 641
- pragmatics, 116, 121–122, 498–499, 650
- propositional contents, 77–78, 113, 135, 228, 401, 485, 586, 649
- reliability inference, 218–219, 221
- representation, 138, 280, 431, 608, 665n31
- semantics, xx, 116, 121–122, 137–138, 650
- substitutional commitments, 115, 319, 417, 467, 512–513
- token recurrence, 451–452
- exportation, 502, 516, 598, 610–611
- expressions, 75–77, 105–106, 384, 392, 393, 399, 403, 591, 650. *See also* sentence(s); singular terms; subsentential expressions; tokenings
- expressive commitments, 516, 545, 586, 588, 589, 608. *See also de re* ascriptions; explicit
- expressive completeness (equilibrium), 111–116, 138, 613, 641–643, 650
- expressive deduction, xxiii, 401, 403. *See also* singular terms
- expressive development, 642
- expressive rationality, 105–111, 116, 125, 130–132, 642
- expressive role, 228, 245–246, 310, 330, 396, 414, 458, 473, 474, 486, 498–499, 541, 590. *See also* anaphora, expressive role of; explicit
- of ascriptions, 502, 504, 505, 522, 529–530, 533, 591, 613
- of conditionals, 108–110, 126–127, 474 of *T*, 554–559, 566
- of identity locutions, 115, 382, 476
- of language, 342, 352, 377
- of logical vocabulary, 95–96, 107–108, 110, 113, 114, 125, 359, 382, 383, 385, 530, 619, 629, 635, 641, 644, 650
- of normative vocabulary, 245–252, 261, 267, 271, 625
- of ‘of’, 138, 704n20
- of representational locutions, 138, 284–285, 330, 499, 502, 505, 522, 529–530, 584
- of semantic vocabulary, 285, 311, 325, 414
- of token recurrence, 310, 453, 454, 458, 473, 474, 486, 590
- of ‘true’, 278, 284–285, 324, 326–333, 568
- expressivism, 92–93, 682n14
- extensionality, 344, 350, 352, 359, 392, 484–485, 668n72, 690n37. *See also* logic, multivalued; substitution
- extensions, 109, 484–485, 681n9, 694n31. *See also* intensions
- externalism. *See* semantic externalism; justification
- facades. *See* barn facades
- facts, xxi, 76, 84, 245, 324, 328, 331, 333, 622, 631–632. *See also* objectivity; objects; representation
- and norms, 137, 331, 623–626, 625
- semantic, 76, 326–329, 331, 333
- as true claims, 327–330, 333, 488, 595, 606, 622, 624, 625, 704n19

- failure, 258–259, 295  
 Feyerabend, Paul K., 480–481  
 Field, Hartry H., 154, 156, 481  
 force, 5, 12, 17, 56, 82, 111, 186–190, 288, 297–299, 322, 339, 343–344, 367, 604, 661n65, 680n60. *See also* assertion(s); content(s); pragmatics; pragmatic significance; speech acts  
 formal inferences, 97–102, 104–105, 107–108, 133–136, 340, 351, 383, 619, 635. *See also* conditionals; material inferences; validity  
 formalism, 97–102, 110, 112, 135, 635  
 formal semantics. *See* semantics, formal/philosophical  
 foundationalism, 90, 177, 204, 216, 221, 681n5. *See also* "Empiricism and the Philosophy of Mind"; 'looks' talk; non-inferential reports  
 frames. *See* substitution frames  
 freedom, 32, 50–51, 662n77, 714n18  
 Frege, Gottlob, 11–13, 23, 80–82, 97, 112, 139, 200, 279, 354, 363, 433, 439, 443, 449, 475–477, 688n29, 693n2. *See also* tactile Fregeanism  
   on concepts, 72, 84, 282, 316, 352, 354, 355, 617, 693n10, 694n34  
   on content, 14, 72, 94–97, 107–108, 123, 281, 288, 298, 340, 344, 345, 347–348, 356, 475, 476  
   and expressive role of logic, 97, 107–111, 113  
   and inferentialism, 11–12, 80, 82, 94–97, 117, 281, 351  
   on objects, 279, 355, 360, 365, 366, 413–415, 419, 421–424, 435, 444, 448  
   on recognition judgments, 417, 418, 442  
   on substitution, 81, 105, 138, 281–282, 346, 367, 369, 414, 422–424, 436–438  
 fruitfulness, 476–477  
 functionalism, 16, 147–149, 159, 160, 196. *See also* incompatibility  
 Gauthier, David, 49  
 'gavagai', 409–412, 429, 430  
 Geach, Peter T., 298, 492–493, 696n51  
 Geach-Frege test. *See* embedding, test  
 Gentzen, Gerhard, 125  
 gerrymandering, 28–29, 36, 41, 62, 208–212, 214, 645, 647, 710n92. *See also* regularism; reliability, reliabilism  
 Gettier, Edmund, 675n1  
 Gibbard, Allen, 682n14  
 giving and asking for reasons, xiv, 20, 46, 54, 136, 200, 209, 212, 230, 232, 275, 278, 330, 403, 449, 453, 496, 498, 590, 628, 637, 648, 715n28. *See also* communication; inferentialism; linguistic practice; reasons  
   and action, 156, 158, 194, 215, 233, 243, 244, 248, 263, 630  
   and assertion, 79, 89, 117, 139, 141, 158, 159, 167, 173, 205, 221, 229, 233, 449, 520, 593, 601–607, 624, 629, 641  
 Goldman, Alvin I., 209–211  
 'good', 289, 298  
 grammar, 304, 361, 404, 406, 688n21  
 grasping, 9, 120, 355, 583, 635, 636. *See also* concepts; tactile Fregeanism; understanding  
 Grice, H. Paul, 146  
 Grotius, Hugo, 18–19  
 Grover, Dorothy, 300–301, 321–322  
 harmony, 124–130  
 Haugeland, John, 34, 36, 37  
 'he', 312  
 Hegel, G. W. F., 50, 85–86, 92–93, 663n1, 669n93, 677n26, 698n78, 716n35  
 Heidegger, Martin, 661n64, 666n35  
 Hempel, Carl G., 679n44  
 'here', 463  
 hierarchy, 36, 51, 160, 216, 242, 362, 404, 408  
   assertional/inferential, 350–353, 355, 356, 359  
   substitutional, 351, 353, 358  
 Hintikka, J., 681n7, 699n85  
 Hobbes, Thomas, 49, 51  
 holism, 89, 92, 426, 477–481, 587  
 Hume, David, 10–11  
 'I', 439, 537, 552–566, 704n23, 707n51, 711n103  
 idealism, linguistic, 331  
 identity, xix, 314, 319, 372, 383, 398, 416, 439, 498, 530, 571, 573, 589, 695n34. *See also* singular terms; substitution claims, 112, 315–317, 324, 418–422, 424, 432, 441, 443, 444, 468, 476, 477, 489  
   locutions, 115, 372, 374, 382, 416–419, 451  
 ignorance, 240, 602, 605. *See also* error imperatives. *See* commands

- implicit. *See* explicit; norms; practice(s)
- incommensurability, 480–481, 483
- incompatibility, 189–190, 225, 332, 429, 560, 602, 634–635, 672n19, 690n35, 691n40, 711n102. *See also* deontic scorekeeping; entitlement; functionalism
- asymmetric/symmetric, 673n23, 674n39, 711n98
- and entailment, 160, 382, 602–603
- and inference, 12, 89, 115, 132, 169, 178, 190–191, 196, 201, 674n43, 710n91, 714n12
- material, 92, 160, 169, 384
- and negation, xix, 115, 436
- and practical commitments, 237, 253, 259, 269–271
- incorrigibility, 292–295. *See also* ‘looks’ talk; trying
- indefinite descriptions. *See* definite descriptions
- indeterminacy of translation, 409–412
- indexicals, 132, 282, 303, 309, 462, 535, 550–553, 558, 560–561, 565, 566, 577–578, 617, 633. *See also* deixis; demonstratives; quasi-indexicals; singular terms; tokenings
- and anaphora, 309, 460, 473, 585
- indirect descriptions. *See* anaphorically indirect definite descriptions
- indirect discourse ascriptions. *See* ascriptions, indirect discourse
- induction, 168, 189–190. *See also* inference(s), entitlement-preserving
- inference(s), 5, 12, 87, 90–91, 97–102, 104–108, 131–134, 139, 189–190, 206, 214–221, 225, 228, 260, 340, 347–348, 351, 372, 377, 383, 385–386, 392, 400, 402, 472, 619, 627, 634–635, 673n28, 689n32. *See also* concepts; content(s); deontic scorekeeping; formal inferences; giving and asking for reasons; inferentialism; linguistic practice; material inferences; representation and assertion, 91, 95, 158, 167–168, 190, 194, 218, 266, 367, 687n11
- commitment-preserving, 168, 189–190, 200, 237, 238, 344, 541, 553, 673n27, 679n40, 689n31, 710n91, 714n12
- committive (*see* inference(s), commitment-preserving)
- and conditionals, 91, 98, 100, 101, 106, 666n48
- and content, 89–92, 95–96, 102–103, 112, 144, 190, 354, 601–607, 618–623, 634–635
- entitlement-preserving, 168, 200, 220, 237, 238, 541, 673n27, 675n44, 678n38, 679n44, 681n2, 710n91, 714n12 (*see also* induction)
- four kinds of, 189–190
- and incompatibility, 89, 132, 189–190, 710n91, 714n12
- inferential strengthening/weakening, 379–382
- permissive (*see* inference(s), entitlement-preserving)
- reliability, 189–190, 215–218, 221, 228
- and representation, xvi, 93–94, 136, 391, 665n31
- social dimension of, 54, 91–93, 138, 158, 197, 358, 518, 519, 593, 601, 605
- substitution, 370–374, 410, 430, 621–622, 689n31, 690n33
- and truth, xvii, 5–6, 96–97, 104–105, 107–108, 277, 689n31
- inference, substitution, and anaphora (ISA), xvi, 198, 281–283, 391, 449–450, 457, 467, 472–473, 495, 621–623, 649, 650. *See also* content(s); deontic scorekeeping; inferentialism; semantics
- inferential articulation, 168–172, 275, 414, 430, 431, 466
- of discursive practice, 79, 91, 132, 142, 157, 186–190, 198, 225, 233–234, 237, 277, 284, 430, 449
- social dimension of, 167, 477–478, 586
- inferential commitments, 116–117, 248, 347, 351, 357, 451, 454, 506, 586–587, 640, 650, 687n11. *See also* conditionals
- inferential involvements, one-way, 371–372, 377, 386, 388, 392. *See also* singular terms; substitution
- inferentialism, xxi, xxii, 93–94, 104, 107–110, 117–132, 135, 137, 200, 214–221, 281, 334–338, 413, 429, 475, 495, 608, 620, 669n90, 686n1, 690n37, 690n37. *See also* concepts; Dummett, Michael; Frege, Gottlob; representationalism; Sellars, Wilfrid; semantics
- and objectivity, 109, 134, 137, 354, 478, 622–623, 633–636
- and representation, xxi, xxii, 92–94, 132, 205, 283, 285, 334–338
- inferentially inverting contexts. *See* contexts, inferentially inverting

- inferential polarity, 381, 387, 398, 399. *See also* contexts, inferentially inverting
- inferential roles, 89–90, 96, 105–106, 114, 281, 349, 413, 429, 618, 620, 636, 667n58. *See also* content(s); grasping; inference(s); inferentialism
- inferential significance, 475, 478, 480–481, 483, 633, 635
- information, 474, 510, 514, 517, 546. *See also* communication; *de re* ascriptions 'ing'/'ed', 330
- inheritance, 168, 169, 283, 306, 460, 472. *See also* anaphora, inheritance of  
of substitutional commitments, 132, 283, 454–455, 472–473, 499, 564, 581, 583, 621–622
- inheritance of entitlements, 168–171, 175–176, 179, 193, 204, 212, 217, 218, 239, 242, 249, 276. *See also* authority; deference; entitlement, default-and-challenge structure of; justification; testimony
- interpersonal, intracontent/intrapersonal, intercontent, 169–170, 175–176, 179, 205, 218, 221, 226, 241–242
- instituting, 165–167, 202, 284. *See also* practice(s); pragmatism  
of norms, xiii, 46–50, 52–55, 137, 140, 498, 626–628, 646  
of statuses by attitudes, 61, 64, 115, 133–134, 142, 161–162, 169, 593, 597, 623–624, 630, 638
- intellectualism, 20–22, 32, 77, 110, 135, 231, 669n92. *See also* Kant, Immanuel; platonism; regulism; rules
- intensions, 482–485. *See also* extensions
- intention(s), 8, 13–14, 58, 146–147, 193, 232, 233, 239, 253–259, 261–267, 523–526, 670n3, 681n1, 696n46, 702n14. *See also* action(s); reliable differential responsive dispositions  
in action, 256–259, 670n6, 678n34, 707n49  
prior, 256–259, 558, 670n6, 678n34, 707n49  
pure, 256–257  
and reasons, 255, 261, 678n37
- intentional explanation, 56, 57, 268–269, 364, 521, 524, 711n95. *See also* action(s); interpretation  
and *de dicto/de re* ascriptions, 522–526  
as normative, 15–18, 195, 268–269
- intentional interpretation. *See* interpretation
- intentionality, xv, 6–7, 15, 59, 61, 70, 99, 148–149, 155, 415, 416, 631. *See also* belief(s); deontic scorekeeping; interpretation; propositional contents; stance(s)  
analogical/relational linguistic approaches to, 16, 150–152  
derivative/original, 60, 61, 143, 171, 629–644, 671n8, 715n25  
descriptive/normative conceptions of, 9, 15, 22, 56–57, 60, 671n8  
discursive, 7, 8, 24, 61–62, 67–70, 142, 631, 649  
linguistic/pragmatic theories of, 16, 22, 76, 148–150, 152, 230, 631  
and representation, 67–70, 336, 547–548  
simple, 59, 171, 631
- intentional stance, 55–62, 67, 629–631, 636–639, 642, 715n27  
simple, 629–630, 639, 642, 643
- intentional states, 16–18, 75, 118, 133, 147, 157, 196, 270. *See also* normative attitudes; normative status(es)  
normative significance of, 8, 13–16, 23, 27–29, 46, 55–57, 62, 67
- intentional systems, 59, 60, 61, 629–630, 642
- interlocutors, 559–560. *See also* deontic scorekeeping
- internalism, 215, 219, 221. *See also* reliability, reliabilism; semantic externalism
- interpretation, 66, 74, 139, 152, 510–513, 628–632, 637, 638, 642, 644–648, 699n86. *See also* deontic scorekeeping; *de re* ascriptions  
and communication, 513, 588, 645, 646, 670n6  
and demonstratives, 510–513  
and deontic scorekeeping, 136, 475, 508, 636, 644–645  
external and internal, 645–648, 715n27, 716n37  
intentional, 55–62, 83, 84, 158, 232, 715n27  
and original intentionality, 632, 640  
and personal pronouns, 510–513  
as substitution, 20–21, 65, 353, 508–513, 591  
and understanding, 508, 517, 658n39, 660n56  
and Wittgenstein, 20–22, 61, 62, 509–513
- interpretive equilibrium, 641–644

- ISA. *See* inference, substitution, and anaphora
- iteration, 295, 313–316, 319, 701n6, 702n13. *See also* anaphora
- I-thou/I-we* sociality, 39, 62, 508, 522, 526, 590, 593, 598–607, 659n50, 716n36. *See also* deontic scorekeeping; norms; objectivity; perspectives; we
- James, William, 287–288
- JTB account of knowledge. *See* knowledge, JTB account of
- judgment, 7, 12, 84–86, 95, 200, 363, 614. *See also* assertion(s); Frege, Gottlob; Kant, Immanuel
- justification, 11–14, 22, 90, 167, 201–202, 204, 217, 221, 228, 294, 515, 532–533. *See also* assertion(s); inference(s); knowledge
- and entitlement, 174, 204–206, 234
- externalism and internalism about, 219, 221
- and reliability, 207–208
- justified-true-belief account of knowledge. *See* knowledge, JTB account of
- Kant, Immanuel, 9–11, 14, 47, 58, 80, 92–94, 102, 200, 230, 337, 475–477, 614, 617, 625, 655n1, 662n93, 665n29, 669n93, 712n10, 713n10
- on acting according to conceptions of rules, 30–33, 35, 41–42, 45, 50, 65
- and classificatory conception of concepts, 85–86, 614–617
- on intentionality as normative, 7–11, 23, 29–30, 289
- on norms as rules, 8, 18–19, 27, 30, 32, 52, 200, 206, 623–624, 655n1, 656n19, 657n31, 712n5
- on primacy of judgment, 79–80, 95, 136, 362–363, 614, 619, 643, 658n41
- three dualisms of, 18, 614–618, 622, 661n65
- on will, 32, 50–52, 270, 271, 665n33
- Kaplan, David, 469, 547–548, 550, 561–562, 566, 617
- knowing-how/knowing-that, 23, 25–26, 135–136, 591, 641, 658n40. *See also* explicit; pragmatism
- knowledge, 74, 177, 200, 202, 204, 209, 210, 221, 228, 297, 645, 676nn3,7, 715n27. *See also* assertion(s); observational knowledge
- and assertion, 199–204
- and doubt, 177, 209
- as hybrid deontic status, 201–205, 213–215, 220, 297
- JTB account of, 201, 207, 228, 297, 515, 675n1
- and reliability, 209–211, 219–220
- and truth, 202, 204, 221, 297
- and understanding, 90, 209, 213–215
- Kremer, Michael, 668n82, 673n20
- Kremer, Philip, 674n42
- Kripke, Saul A., 28, 37, 209, 322, 468–471, 483, 488, 575–579, 582, 603, 656n10, 660n52, 713n11
- Kulas, J., 699n85
- Kvart, Igal, 691n46
- Lakatos, Imre, 697n66
- Lance, Mark, 668n73, 673n23, 674n42, 690n35
- language(s), 24, 146–147, 232, 342, 352, 365, 377, 403, 411, 500, 650. *See also* assertion(s); deontic scorekeeping; explicit; giving and asking for reasons; inference(s); linguistic practice; norms; we
- and *de re* ascriptions, 499–503
- entry/exit transitions, 335–336, 528, 632. *See also* action(s), as language
- exit transition; perception, as language entry transition
- games, 91, 172, 179
- and mind, xv, xxiii
- natural, 145, 411, 499, 504, 520, 688n21
- use (*see* pragmatics)
- Leibniz, Gottfried Wilhelm, 10, 93, 337, 614
- LePore, Ernest, 690n32
- Lewis, C. I., 616
- Lewis, David K., 147, 180–183, 185, 187, 232, 404, 460, 483, 550, 670n3
- licenses, 161, 163–165, 196
- Lindenbaum algebras, 342, 345
- linguistic practice, xi, 16, 141, 155–156, 167–168, 172, 232, 275–277, 331–332, 360, 496, 586, 623, 628, 630–632. *See also* assertion(s); deontic scorekeeping; explicit; giving and asking for reasons; inference(s); norms

- and assertion, 141, 167–168, 172, 173, 276, 367, 586, 623, 628–629, 686n1  
 inferential articulation of, 141–143, 156–158, 167–168, 183, 279–281, 431, 449, 608, 614, 637, 649
- Locke, John, 146, 614
- logic, 12, 96–100, 108–111, 110, 135, 231, 340, 347–348, 353, 384, 435. *See also* explicit; logical vocabulary  
 expressive approach to, xix, 108–111, 117, 125, 131, 650  
 multivalued, 340–346, 358, 359
- logical vocabulary, xviii, xx, 76, 101, 103–105, 112, 114, 116, 123–125, 127, 135, 266, 346, 381–384, 394, 414, 498, 530, 559, 566–567, 601–607, 639–643, 667n67, 691n42. *See also* ascription(s); conditionals; deontic scorekeeping; explicit; identity; negation; normative vocabulary; quantification  
 expressive role of, 13, 95–96, 109, 113, 116–117, 131, 335, 350, 359, 382–383, 385, 392, 402–403, 530, 619, 629, 635, 641, 643–644, 650  
 and singular terms, 393, 395–397, 401  
 ‘looks’ talk, 292–297, 681n5. *See also* foundationalism
- Lycan, William, 691n47
- making-true, 5, 8, 13, 46, 233, 236, 277, 521, 681n1. *See also* action(s); phenomenalism, normative; practical commitments; taking-true
- Manders, Kenneth, 691n41
- maps, 65, 74, 518, 665n31. *See also* representation
- material inferences, 97–102, 104–105, 125–137, 189, 206, 218, 345, 349, 359, 373, 374, 383, 402, 619, 666n48, 690n36. *See also* conditionals; content(s); enthymemes; explicit; formalism; incompatibility; inference(s); inferentialism; Sellars, Wilfrid  
 and conceptual contents, 98, 102–108, 618, 623, 634–635  
 and formal inferences, 97–102, 104–105, 117, 135–136, 345, 383, 666n41, 672n19
- McDowell, John H., 29, 54, 318, 535–539, 561–562, 566, 567, 577, 582, 603, 655n2, 658n39, 660n59, 686n58, 689n29, 693n10, 705n26, 707n55, 708n72
- McGinn, Colin, 703n17
- meaning, xii–xiii, 4, 13–14, 60, 62, 73, 88–89, 121–124, 126, 146, 478, 649, 716n32. *See also* content(s); inference, substitution, and anaphora; inferentialism; norms; pragmatics; semantics; understanding
- Meinong, A. Alexius, 71
- methodology, 229–233, 526, 592. *See also* explanatory strategy  
 explicit/implicit, xiv, 109, 587  
 phenomenalism, 597, 636–637  
 substitutional, 81, 95–96, 104–105, 138, 281
- mind, xv, xxiii, 650. *See also* explicit; inference(s); intentionality; language(s); rationality; sapience; sentience; we mistakes, 21, 27–28, 31, 52, 54, 258
- modality, 13, 105–106, 318, 690n36, 703n19
- model theory, 667n58
- Moore, G. E., 299
- multivalued logic. *See* logic, multivalued
- multivalued, 340–346, 348–353, 358. *See also* designatedness; logic, multivalued; substitution
- ‘must’, 12, 14, 30
- name-bearer model, 69, 84, 352, 359. *See also* designation; nominalism; proper names
- names. *See* proper names
- naturalism, xiii, xv, 10–15, 31, 42–46, 289, 299–300, 460–461. *See also* causes; dispositions; norms; regularism; regularities; reliability
- natural kind terms. *See* sortals
- Neale, Stephen, 491–494
- necessity, 9–11, 30, 624. *See also* Kant, Immanuel
- negation, xix, 92, 115, 319, 381–382, 393, 436, 498, 530, 692n54. *See also* explicit; incompatibility; logical vocabulary  
 expressive role of, 401, 498
- negative existential claims. *See* existential commitments, negative existential claims
- nominalism, 84, 665n32. *See also* designation
- noninferential reports, 64–66, 88–91, 207, 209–210, 212, 214–221, 225, 236, 293, 335–336, 458, 473, 531–532, 620,

- noninferential reports (*continued*)  
 674nn38,40. *See also* content(s), empirical (cognitive); foundationalism; inferentialism; observational knowledge; perception; reliable differential responsive dispositions  
 authority of, 122, 212, 215–217, 221, 223–227, 234  
 as not autonomous, 216, 465–466  
 and circumstances and consequences of application, 119–120, 221, 226  
 enabling/defeating conditions, 226–227  
 Sellars on, 215–217  
 and solidity of linguistic practice, 631–632
- normative attitudes, 32–37, 39, 42, 45, 47–50, 52, 54, 64, 626, 627, 639. *See also* assessment; deontic attitudes; normative status(es); norms; pragmatics, normative; sanctions
- normative phenomenalism. *See* phenomenalism, normative
- normative pragmatics. *See* pragmatics, normative
- normative significance, 48–50, 52, 54, 167, 656n17. *See also* instituting; norms; pragmatic significance  
 of intentional states, 8, 13–16, 55–57, 62, 67
- normative stances, all the way down, 638
- normative status(es), 16–18, 33, 37, 39–41, 47, 64, 627–628, 676n4. *See also* deontic status(es); instituting; norms; practice(s); pragmatics, normative
- normative vocabulary, xiii–xv, xviii, 47, 116, 233, 246–249, 267, 624, 625, 637, 640. *See also* explicit; norms; practical reasoning  
 expressive role of, 250–252, 267, 625
- norms, xx, 7–18, 20–21, 30, 35, 37–42, 44–46, 49–52, 55–62, 200, 226, 636–639, 641, 648, 649, 656n18, 661n69. *See also* deontic scorekeeping; explicit; instituting; practice(s); pragmatics  
 all the way down, 44, 625, 638, 649, 660n59, 714n20  
 and causes, 27–29, 33, 45–50, 625, 626, 661n65, 687n2  
 conceptual, 12–14, 46, 53–55, 624, 636, 637  
 and facts, 623–626  
 implicit in practices, xiv, 20–22, 25, 26, 29–30, 41, 45, 46, 54, 55, 62, 623, 626, 627, 631, 648, 649  
 objective, xxiii, 52–55, 63, 253, 631  
 origin of, 626–628  
 pragmatism about, 21–23, 32  
 regularism/regulism about, 18–23, 26–29, 41, 46, 99, 648  
 as rules, 7–11, 18–23, 32, 200, 624, 656n19
- novel terms, 420, 421, 423, 435, 442
- 'now', 463, 559
- numbers/numerals, 437, 442–444, 449
- 'object' as pseudosortal, 438
- object-dependence, 567–573. *See also* beliefs, strong *de re*; existential commitments; singular thoughts
- object-directedness, 415, 416
- objectivity, xvii, 78, 136–138, 253, 498, 529, 592–597, 602–607, 649, 672n18. *See also* I-thou/I-we sociality; norms; perspectives; representation  
 of conceptual norms, xxiii, 53–55, 63, 497, 529, 593–594, 599–607, 631, 633, 636  
 and deontic attitudes/statuses, 599–601, 627  
 and *de re* ascriptions, 598, 600–601  
 as form of contents, 597, 600–601  
 and intersubjectivity, 599–607  
 perspectival character, 52–55, 197–198, 592–602, 604, 609  
 and representation, 140, 280, 530, 609, 672n18, 677n11
- objectivity proofs, 601–607
- objects, 292, 333, 360, 403, 438, 571, 649, 694n29, 713n10. *See also* expressive deduction  
 abstract, 84, 421, 422, 449, 695n34  
 and facts, 84, 331, 333, 622  
 picking out, 375, 413–417, 423, 425, 430–432, 444, 451, 462  
 recognizing as the same again, 416–419, 424, 425, 467  
 representation of, 136, 280, 333, 337, 617, 665n29  
 and singular terms, xxi, 69–70, 414–416, 687n14  
 and solidity of practice, 631–632  
 why are there?, xxii, 404

- observational knowledge, 209, 216–226, 297. *See also* noninferential reports
- observation reports. *See* noninferential reports
- occurrence, 342, 373–375, 389, 454, 465. *See also* recurrence; substitution
- primary, 374–375, 378–379, 381, 389, 392, 394, 395, 397, 400, 421
- 'of', 68, 138, 501–503, 505–506, 516, 520, 543, 546, 565, 569–570, 585–586, 590, 598, 608, 704n20. *See also de re* ascriptions
- ofness. *See de re* ascriptions; 'of'
- 'one', 438
- ontological commitments. *See* existential commitments
- opacity, 571, 574–575, 583
- 'ought', 5, 31, 56, 61, 252–254, 270, 271, 289, 290. *See also* intentional explanation; practical reasoning
- paradoxes, semantic, 321–322
- paratactic theory of ascriptions. *See* ascriptions, paratactic theory of
- parrots, 88, 122
- particularity, 620, 623, 687n14. *See also* deixis; inference, substitution, and anaphora; objects
- paycheck sentences, 490–492
- Peirce, C. S., 289
- perception, xv, 7–8, 119–120, 122, 131, 142, 209–211, 261, 276, 332, 556, 618, 704n22. *See also* content(s), empirical (cognitive); deontic scorekeeping; noninferential reports; observational knowledge; reliability; reliable differential responsive dispositions
- and action, 233–238
- as language entry transition, 221, 233–235
- performances. *See* action(s); speech acts
- Perloff, Michael, 702n14
- permissive inferences. *See* inference(s), entitlement-preserving
- Perry, John, 460, 550, 552–553, 558, 561, 566
- perspectives, 37, 62, 197, 503, 591, 594–595, 598, 608, 630, 635, 710n95, 715n27. *See also* deontic scorekeeping; *de re* ascriptions; *I-thou/I-we* sociality; objectivity; representation
- and ascriptions, 508, 549, 584, 590, 591, 613
- and attributing/undertaking, 61, 508, 597, 598, 649, 677n11
- and communication, 139, 635, 647
- and contents, 139–140, 485, 517, 529, 586–597, 601, 635, 636
- and knowledge, 202, 205, 297, 715n27
- privileged, 599–600, 604, 606
- and semantic externalism, 633, 647
- phenomenalism, 291–297, 327, 624, 631, 682n14, 709n75. *See also* deontic scorekeeping; 'looks' talk; norms; objectivity; stance(s)
- about deontic statuses, 334, 339, 637
- about norms, 25, 280, 627
- about truth, 287, 291, 292, 296–297, 322
- generic, 295–296
- normative, 627, 636, 637, 644
- and pragmatism, 296, 322
- subjective, 292–297
- philosophical semantics. *See* semantics, formal/philosophical
- physicalism, 47. *See also* causes; norms; regularism
- Plato, 201
- platonism, 20–22, 110, 231. *See also* intellectualism
- points of view. *See* perspectives
- practical commitments, 233, 237, 244, 245, 253–256, 262–266, 276, 679n40. *See also* action(s); content(s), practical; deontic scorekeeping; norms; practical reasoning
- causal efficacy of, 259–262, 271, 596
- and doxastic commitments, 233, 236, 238–243, 271
- entitlement to, 238, 253–256, 265–266
- inferential articulation of, 233–234, 237
- and intentions, 193, 256–259, 525, 702n14
- practical reasoning, 7, 83, 100, 116, 158, 233, 240, 244, 246–253, 268–269, 287, 290, 521, 537, 640. *See also* action(s); deontic scorekeeping; explicit; intentional explanation; material inferences; norms
- and action, 8, 231, 336, 507
- and doxastic and practical commitments, 237, 243, 246, 253, 254



- practical reasoning (*continued*)  
 and entitlement-preserving inferences, 249–250  
 and normative vocabulary, 233, 247–249, 625, 640  
 patterns of, 233, 245–252, 271  
 and ‘shall’/‘should’, 245, 258–259, 261, 263–264, 267–271, 553, 680n56, 60  
 and theoretical reasoning, 520–522  
 and will, 233, 270
- practice(s), xi, 7, 22, 25, 41, 52–55, 61, 62, 65, 83, 91–93, 130, 159, 165–168, 205, 233, 339, 346, 627–628, 638, 660n51. *See also* conferring content; deontic scorekeeping; explicit; giving and asking for reasons; instituting; linguistic practice; norms; regularism; regulism and contents, 64, 77–79, 145–147, 497  
 idealized, 128, 168  
 norms implicit in, 25, 29–30, 55, 625, 628, 648, 649  
 and regularities, 26, 625, 638  
 and rules, 20–22, 32, 66, 91, 99–100, 110  
 solidity of, 332, 528, 631–632, 686n55, 715n28
- pragmatic priority of the propositional, 79–81, 83, 95, 337. *See also* assertion(s); Frege, Gottlob; Kant, Immanuel; Wittgenstein, Ludwig
- pragmatics, 133, 140, 159, 187, 212, 334, 474, 592, 624, 681n6. *See also* circumstances and consequences of application; deontic scorekeeping; force; pragmatic significance; speech acts  
 normative, xiii, 132–134, 140, 199, 623, 649  
 and semantics, xiii–xv, xvii, xxii, 64, 68, 83–85, 91, 132–134, 140, 143, 363, 496, 498, 649, 686n1, 704n20
- pragmatic significance, 118, 142, 157, 163, 168, 173, 182–186, 262, 339, 343, 345, 359, 386, 392, 424, 450, 710n91. *See also* content(s); deontic scorekeeping; instituting; norms; practice(s)  
 of assertions, 157, 168, 190, 200, 234, 358, 637
- pragmatism, xii, 21–23, 76, 101, 110, 120, 123, 132, 137, 143, 205, 289, 296, 300, 322, 496, 591–592, 686n1. *See also* norms; practice(s); regulism; Wittgenstein, Ludwig  
 about norms, 21–23, 42, 55, 509  
 classical, 123, 285–291, 296, 303, 322  
 and truth, 285–286, 288–291, 297  
 predicates, xv, 320, 369, 378, 380, 391–395, 398, 400, 401, 406, 410, 411, 617, 620, 622, 623. *See also* singular terms; substitution  
 complex/simple, 371, 406, 407, 409, 434, 436  
 preferences, 240, 246, 248, 678n29. *See also* practical reasoning  
 prima facie reasons. *See* reasons, prima facie  
 Prior, Arthur N., 125  
 prior intentions. *See* intention(s), prior  
 privilege, 28–29, 36, 43, 63, 104–105, 117, 292–293, 634–635. *See also* perspectives, privileged  
 pro-attitudes, 246–247, 267. *See also* normative vocabulary; practical reasoning  
 proform-forming operators, 283, 313–315, 318, 325  
 projection, 366, 395, 396, 399, 636, 647, 688n23  
 promises, 163–165, 172, 180, 262–266, 289, 672n20. *See also* practical commitments  
 pronoun-forming operators, 305, 306, 322  
 pronouns, 132, 301, 305, 308–309, 311–331, 450, 456, 458, 460, 469, 486, 491, 510–513, 582, 707n51. *See also* anaphora; ‘refers’; singular terms  
 proof theory, 667n58  
 proper names, 308–309, 420, 439, 469, 470, 549, 573–579, 581–583. *See also* singular terms; substitution; token-recurrence structures  
 anaphoric account of, 308–309, 458, 459, 470, 572–574, 579–583, 585–586  
 causal-historical theories of, 470–471, 580, 713n11  
 properties, 9, 12, 47, 52, 324, 697n67. *See also* proprietries  
 propositional, pragmatic priority of. *See* pragmatic priority of the propositional  
 propositional-attitude-ascribing locutions. *See* ascription(s), of propositional attitudes  
 propositional attitudes, ascriptions of. *See* ascription(s), of propositional attitudes  
 propositional contents, xiv–xv, 5, 8, 11, 75, 77–78, 83, 135, 141, 168, 203, 209, 221, 230, 275, 329, 333, 335, 339, 413,

- 495, 520, 567–568, 624, 640, 649, 665n32, 702n14, 705n32. *See also* asser-tion(s); content(s); explicit; giving and asking for reasons; linguistic practice; objectivity; objects; representation and assertion, xxiii, 12–14, 157, 187, 221, 278, 335, 586, 604
- conferring of, 63, 275, 277, 279, 284, 623, 645
- de re* specifications of, 513–517, 520
- and inference, 6, 91–92, 101, 104, 200, 209, 277, 281, 337–339, 413, 449, 495
- perspectival character of, 485, 497, 630, 636
- representational dimension of, 75, 84, 153, 333, 414, 474, 485, 508, 519–520, 583, 649
- and truth conditions, 6, 277, 278, 326, 329
- proprieties, 9, 12, 25, 47, 52, 63, 83, 159, 411, 628, 636, 676n4, 677n11, 688n23. *See also* norms; practice(s); properties
- prosentence-forming operators, 303–305, 327
- prosentences, 283, 301–306
- prosentential theory of truth. *See* truth, prosentential theory of
- prosortals and pseudosortals, 438–439
- psychologism, 11–12, 23
- Pufendorf, Samuel, 18–19, 46–50, 160, 662n75
- punishment, 34–36, 40, 42, 43, 63, 179, 662n75. *See also* sanctions
- purport. *See* representation, purport/up-take; singular referential purport
- Putnam, Hilary, 299, 481
- “Puzzle about Belief,” 573–579
- quantification, 300–301, 303–304, 307, 321, 323, 382, 435–437, 439, 440, 491–494, 530, 696n51, 711n100. *See also* explicit; logical vocabulary; predi-cates; substitution
- expressive role of, 383, 396, 399, 434–436, 498, 696n50
- quasi-indexicals, 559–567, 569–570, 581, 588, 590, 707n63, 711n103. *See also* ascription(s); communication; explicit
- quasi-indicators, 563–564, 566
- queries, 172, 192–193. *See also* deontic scorekeeping; speech acts
- Quine, W. V. O., 92, 129, 223, 296, 342, 360–361, 409–412, 477–482, 547–548, 550, 587, 634, 704n24
- quotation
- < > type designation, 314
  - // token designation, 314
  - ⌈ ⌋ corner quotes, 688n25, 704–705n24
- and direct/indirect discourse ascriptions, 529–534, 537
- <sup>s s</sup> scare quotes, 545–547, 588–590, 672n19
- ralugnis mrets, 390–391
- rational agency. *See* action(s); practical rea-soning
- rationalism, 10, 85–87, 92–94, 102, 103, 614. *See also* inferentialism
- rationality, 5, 6, 15, 17–18, 50–51, 56–57, 62, 99, 108–110, 127–128, 131, 135, 183, 195, 196, 229–233, 244–245, 252–253, 271, 364, 643. *See also* explicit; giving and asking for reasons; inten-tionality; interpretation; linguistic practice
- expressive, 105–106, 116, 125, 130–132
- instrumental, 106, 108–110, 131, 231
- and will, 32, 50–51, 661n69, 665n33
- reasons, xxiii, 8, 11–12, 116, 141, 171, 200, 203, 253, 275, 276, 383, 402, 449, 474, 590, 618, 649. *See also* giving and ask-ing for reasons; inference(s); inferen-tialism; norms; practical reasoning; rationality; theoretical reasoning
- and action, 56, 171, 230, 237, 240, 244, 245, 249–253, 261, 678n37
- and causes, 234, 259–260, 271, 661n65
- force of, 5, 17, 56, 354
- prima facie, 249–253
- primary, 240, 246, 259, 262, 663n84
- recognition, 4, 31–32, 67, 640. *See also* we
- recognition judgments, 112, 417–419, 424, 432, 444. *See also* Frege, Gottlob; iden-tity, claims; singular terms; substitu-tion
- recurrence, 455, 457–459, 463, 465, 468. *See also* anaphora; occurrence; substi-tution; token-recurrence structures; tokenings
- recurrence commitments, 456–457, 472
- redundancy, 288–289, 291, 299–300, 303–304, 316, 322, 329. *See also* truth

- reference, 84, 305–327, 355, 356, 360, 467, 471, 477–482, 550, 703n16. *See also* explicit; objects; ‘refers’; representation; singular terms; substitution; truth
- and anaphora, 304, 306, 307, 325
- direct, 464, 550, 567–568, 571–574, 582
- and inference, xvi, 136, 391
- as a relation, 306, 323, 325–326
- reference classes, 210–212. *See also* reliability
- referential purport. *See* representation, and purport/uptake; singular referential purport
- ‘refers’, xvii, 116, 137, 138, 279, 283–285, 305–306, 312–322, 325, 464. *See also* anaphora; reference
- and anaphora, 305–306, 323
- and representation, 336, 499
- and ‘true’, 283–285, 305–307
- reflection, Socratic, 105–106, 127–128, 130
- reflexives, 563, 573, 698n80
- regress, 22, 31–32, 36, 61, 74, 77, 177–178, 205, 206, 221, 451–452, 519, 600–601, 646, 647. *See also* pragmatism; regulism; Wittgenstein, Ludwig
- of rules, 20–30, 36, 45–46, 61–66, 206, 509, 657n31
- regularism, 26–29, 32, 36–42, 46, 62, 208, 594, 658n35. *See also* norms; regularities
- regularities, 11, 28–29, 37–42, 624, 629, 659n47, 660n52, 710n92. *See also* regularism
- and correctness, 27–28, 62, 207–208, 212
- and norms, 26–29, 45–46, 99, 628, 648
- and practices, 625, 638
- privileged, 28–29, 36, 63
- and reliability, 207–209, 211
- regulism, 18–29, 32, 62–63, 110, 594. *See also* intellectualism; norms; pragmatism; rules
- Reichenbach, Hans, 683n34
- reinforcement, 34–35, 37, 42–43, 659n45. *See also* dispositions; sanctions
- reliability, 207–214, 226, 331, 427, 532–533, 556, 646, 681n2. *See also* barn facades
- attribution of, 217–221
- and authority, 215–217, 219–220
- and entitlements, 167, 207–208, 210–211
- and inference, 189–190, 215–218, 221, 228
- and regularity, 207–209, 211
- reliabilism, 121, 206–209, 219–220, 228, 715n27
- and success, 521, 527–528
- reliable differential responsive dispositions, 5, 33, 35, 42, 89, 119, 122, 162, 210, 214–221, 223–225, 263, 269–271, 427, 430, 473, 518, 531, 555, 556, 618, 658n45. *See also* perception
- and action, 234, 262, 269–270, 524–525, 702n15
- and inferentialism, 86–88, 91, 119–120, 156, 429, 622
- and noninferential reports, 7–8, 209, 223–224, 235, 261, 293, 465
- repeatables, 282, 449–452, 454, 592, 623. *See also* token-recurrence structures
- replacement, 369, 371, 394, 396, 406–409, 692n51. *See also* substitution frames
- reports. *See* noninferential reports
- representation, xvi, 7, 10–11, 17, 31–32, 69–72, 74, 77–79, 88–90, 116, 136, 330–333, 337, 391, 414, 495, 502, 518, 519, 609, 617, 649. *See also* aboutness; inferentialism; objectivity; propositional contents; representationalism; semantics
- and content, 6, 54, 69, 70–72, 78, 79, 84, 333, 474, 497, 508, 519–520, 590, 593–595, 601–607, 609, 649
- and deontic scorekeeping, xvii, 187, 324, 584
- and *de re* ascriptions, 516, 547–548, 552, 586
- and discursive practice, 279–282
- expression of, 75–77, 335–336, 518–519, 608, 665n32, 709n90
- and inference, 93–94, 518–519, 665n31
- and intentionality, 67–70, 279
- and objectivity, 54, 75, 78, 140, 151, 153, 280, 630, 677n11
- and perspectives, 496, 497, 529
- and purport/uptake, 6–7, 62, 70–75, 77, 78, 89–90, 138, 489
- and success, 70–75
- representational contents, 6, 69, 74, 75, 84, 135–140, 497, 517, 520, 528, 609. *See also* propositional contents

- representationalism, xxii, 6, 31–32, 84, 92–94, 97, 334–338, 669n90. *See also* empiricism; inferentialism
- representational locutions, 279, 284, 306, 496, 500, 519, 584. *See also de re* ascriptions; explicit; 'refers'
- 'represents', 70–72, 517. *See also* representation
- responsibility, 7, 10–11, 17–18, 163, 170, 172, 173, 178, 179, 254, 516, 624, 712–713n10. *See also* assertion(s); justification; norms
- and authority, xii, 162–163, 165, 242, 673n24
- rewards, 34–36, 40, 42, 63. *See also* sanctions
- rigidity, 14, 468–472, 486, 573, 698n80.  
*See also* anaphora
- Rosenberg, Jay F., 147
- Rousseau, Jean-Jacques, 51, 662n70
- Royce, Josiah, 375
- rules, 8, 10–11, 20–22, 25, 31–32, 64–66, 91, 130, 616, 658n40. *See also* intellectualism; norms; pragmatism
- conceptions of, 9, 30–32, 65, 624
- and correctness, 18–19, 23, 24, 26, 64–66
- elimination/introduction, 101, 117–130
- and norms, 7–11, 15, 18–23, 62, 624, 656n19
- and practices, 20–22, 91, 99–100, 110
- regress of, 20–27, 29–30, 36, 45–46, 62–66
- and regulism, 18–23
- Russell, Bertrand, 307, 415, 436, 437, 567–568
- Ryle, Gilbert, 23
- Salmon, Nathan, 708n67
- samesaying, 536–537, 539, 562, 566
- sanctions, 34–37, 42–46, 63, 162, 164, 179, 659n45
- external/internal, 44–46, 162–165, 178–180, 660n54, 662n75
- and normative attitudes, 35–36, 42, 45, 166
- and reinforcement, 34–35, 42–43
- sapience, 4–8, 87, 88, 231, 275–277, 520, 559, 591, 644. *See also* linguistic practice; rationality; sentience; we
- saying, 18–19, 64–66, 107–108, 336, 535–538, 629, 647, 649. *See also* assertion(s); explicit; 'says'
- and doing, 30, 62, 77, 108–110, 115, 135, 639–641, 658n40
- 'we', 3–5, 7, 643–645
- 'says', 529–539, 542. *See also* saying
- Scheffler, Israel, 481
- score. *See* deontic scorekeeping
- scorekeepers. *See* deontic scorekeeping
- Searle, John R., 60, 70, 147, 256–257
- seeming. *See* 'looks' talk
- self-consciousness. *See* consciousness/self-consciousness
- Sellars, Wilfrid, 16, 105–107, 130, 155, 209, 234, 258, 306, 660n51, 679n39, 706n39
- on concepts, 10, 89–91, 93, 94, 96, 117, 484, 618, 634
- on giving and asking for reasons, 139, 167, 263
- on 'looks' talk, 293–294
- on material inferences, 102–103, 105–106, 109, 618
- on noninferential reports, 122, 214–221, 228, 465–466
- on regress of rules, 23–26, 184
- selves, 554, 555, 559–560. *See also* 'I'
- as coreresponsibility classes, 559–560
- semantic externalism, 219, 332, 631–633, 645–649. *See also de re* ascriptions
- semantic facts. *See* facts, semantic
- semantic primitives, 79, 133–134, 136, 283, 285, 681n6, 687n2
- semantics, 140, 144, 187, 360, 363, 370–372, 592, 687n2. *See also* content(s); inferentialism
- agent, 147, 149, 151, 671n8, 696n46
- formal/philosophical, 143–145, 199
- inferential, xiii, 93–94, 132–134, 334–338, 360, 623, 649
- and pragmatics, xiii–xv, xvii, xxii, 64, 68, 83–85, 91, 132–134, 363, 498, 592, 649, 686n1, 704n20
- semantic vocabulary, xvi, 82, 145, 279, 280, 283–285, 325. *See also* explicit
- traditional, 116, 285, 323, 325–327, 330, 432, 450
- semiotics, 80
- senses, 340, 343, 355, 356, 536–537, 570, 571, 574–575, 579, 582–583, 681n9.  
*See also* content(s); tactile Fregeanism
- sentence(s), 5, 69–70, 81, 82, 84, 118, 146, 173, 288, 329, 338, 347, 363, 364, 366, 377, 385, 391–395, 409–412, 426, 453,

- sentence(s) (*continued*)  
 533, 539, 669–670n1, 688n29. *See also* deontic scorekeeping; explicit; giving and asking for reasons; inference(s); linguistic practice; pragmatic priority of the propositional; propositional contents  
 and assertion, 82, 85, 141, 156, 157, 168, 199, 277, 367, 586  
 and beliefs, 146–147  
 compound, 199, 340–342, 345, 347–349, 351, 352, 355, 358, 395, 397, 398, 402, 403  
 frames, 377, 384, 386, 393, 394, 403, 405, 407  
 freestanding, 199, 298–300, 392, 399  
 nominalizations, 303–306, 516  
 novel, 362–367, 375, 414  
 sentence, 5, 7, 275–277. *See also* sapience  
 ‘shall’, 258–259, 261, 267–271, 680nn56,60, 716n34. *See also* practical reasoning  
 and practical commitments, 245, 261, 263–264, 269, 271, 553  
 sharing concepts, 562, 590, 631–636. *See also* communication  
 ‘should’, 267–271, 716n34. *See also* practical reasoning  
 and practical commitments, 269, 553  
 significance. *See* normative significance; pragmatic significance  
 simple material substitution-inferential commitments (SMSICs), 372–376, 378, 382–385, 389–392, 397, 398, 400, 410, 424, 432, 434, 444, 450, 454, 455. *See also* identity; singular terms; substitution; substitutional commitments  
 asymmetric/symmetric, 375, 377, 379–382, 386, 388, 393–396  
 and singular terms, 424, 432, 454  
 and substitution inferences, 374–376  
 singular referential purport, 308, 361, 415, 416, 418, 422, 425, 426, 432, 442, 571. *See also* objects; representation; singular terms  
 success of, 431, 433, 441, 443, 444  
 singular terms, 308, 323, 360–362, 370, 376, 385–391, 398–401, 404, 409–411, 413–419, 422, 426, 432–434, 443, 453, 454, 539, 560, 687n14, 694n17, 695n40  
 in compound sentences, 395–397  
 and content, 198, 396, 425, 453, 617, 623  
 dual substitutional characterization of, 385, 390, 397, 400  
 and logical vocabulary, 379–382, 393, 395  
 maximal/minimal substitutional requirement on, 419–426  
 and objects, xxi, 69–70, 288, 307, 326, 414–416, 578, 687n14  
 primary occurrences of, 342, 374–375, 378–379, 381, 389, 392, 394, 395  
 semantic/syntactic substitutional role of, 363, 376, 378, 380, 384–385  
 and sentences, 81, 82, 669n1, 688n29  
 substitution-inferential significance, 370–372, 376, 378, 400, 403  
 and substitution-structural roles, xv, 324, 369, 371–373, 378, 384–390, 401, 406, 424, 454, 619–620, 622  
 what are?, 361, 367, 399, 400  
 why are there any?, 376, 378–381, 384, 385, 399–401, 403, 414  
 singular thoughts, 567–574, 583  
 SIS. *See* substitution-inferential significance  
 SMSICs. *See* simple material substitution-inferential commitments  
 social dimension, 496, 497, 520, 590, 601–607  
 of inference, 138, 519–520, 586, 593, 605  
 social perspectives. *See* perspectives  
 social practices. *See* practice(s)  
 solidity of practices, 332, 528, 631–632, 686n55, 715n28  
 Solomon, W. David, 680n56  
 sortals, 304, 409–412, 420, 437–440, 445–446, 461, 471, 684nn36,38, 690nn34,39, 696n51. *See also* singular terms  
 Sosa, Ernest, 550–551, 566, 701n2  
 speaker’s reference, 487–490, 511–513, 703n16  
 speech acts, xiii, xiv, 37, 82, 133, 142, 151, 157, 159, 172, 173, 182–186, 194, 230, 263, 343, 363, 367, 370, 531, 630, 636, 637, 649, 675n45. *See also* assertion(s); deontic scorekeeping; pragmatics  
 Spinoza, Benedictus de, 93  
 Sprachspielen. *See* language(s), games  
 SSR. *See* substitution-structural roles  
 Stalnaker, Robert, 148–149, 230, 337, 670n3, 671n12, 677n27

- stance(s), 55–62, 280, 287–289, 639, 644.  
*See also* intentional stance; interpretation; phenomenalism  
 discursive scorekeeping, 628, 630, 639–644, 648  
 and intentionality, 55–62, 639–641
- states of affairs, 69–70, 76, 77, 85
- statuses. *See* deontic status(es); normative status(es)
- stimuli, 427–429. *See also* reliable differential responsive dispositions
- stipulation, 145–147, 671n8. *See also* conferring content; semantics
- STIT, 702n14
- Strawson, Peter F., 288, 369, 371
- strong *de re* ascriptions. *See de re* ascriptions, strong
- strong *de re* beliefs. *See* belief(s), strong *de re*
- structuralism, 80
- subordinate/superior, 241–243
- subsential expressions, 136, 342, 360, 362–367, 375–378, 386, 392, 401, 402, 411, 414, 619–620, 640. *See also* singular terms; substitution
- substituted for, 301–303, 368–372, 376–386, 389, 392–398, 400–401, 403–405, 435, 436. *See also* singular terms; substitution; substitution-structural roles
- substituted in, 368–372, 378, 386, 404, 405, 409. *See also* substitution; substitution-structural roles
- substitution, xv, 65, 80, 112, 281–282, 344, 346–348, 367, 369, 372, 383, 391, 394, 403, 405, 430, 452, 509, 619, 623, 696n50. *See also* inference, substitution, and anaphora; semantics; singular terms  
 and anaphora, 450–452, 465, 621–623  
 and content, 347–349, 354–357  
 and inference, 198, 283, 347–348, 352, 370–372, 414, 422, 425, 426  
 intersubstitution, 308–310, 314, 324  
 invariance under substitution, 281, 309–310, 317, 452  
 and multivalued logic, 341–346, 352, 358, 359  
 and subsential expressions, 346–350, 359, 367, 399  
 substitutional invariance, 281, 309–310, 317, 368  
 substitutional variance, 367–369, 371, 379–381, 395, 405, 423, 451, 455  
 substitutional commitments, xxii, 138, 282, 283, 319, 324, 325, 409, 419, 422, 426, 431, 433–440, 447, 451, 454, 459, 467, 472–474, 495, 499, 506, 640. *See also* explicit; inference, substitution, and anaphora; simple material substitution-inferential commitments; singular terms  
 substitution frames, 368–369, 371, 376–381, 386, 394–396, 400–401, 403, 404, 406, 407, 692nn51,54. *See also* predicates; substitution-structural roles  
 substitution inferences, xv, 282, 319, 370–372, 374, 379–381, 450, 477, 619–622, 690n33. *See also* simple material substitution-inferential commitments  
 substitution-inferential significance, 370–372, 376, 378, 400, 403  
 asymmetric/symmetric, 377, 378, 384–385, 388, 391–395, 403  
 substitution licenses. *See* identity, locutions  
 substitution-structural roles, 367–371, 376, 378, 384–385, 400, 403, 404
- success, 13, 203, 236, 243, 280, 287, 290–291, 433, 521, 524, 526–528, 646
- supervenience, 47, 52, 292, 295–296, 628
- tactile Fregeanism, 579–583, 635. *See also* anaphora; content(s); senses
- taking-true, 5, 8, 11, 13, 46, 82, 202, 203, 228, 231, 233, 236, 240, 277, 278, 287–291, 297–299, 322, 521, 599, 625, 627. *See also* assertion(s); making-true; phenomenalism
- Tarski, Alfred, 84, 279, 317, 361, 364
- Tarskian contexts, 316, 318
- Tarski biconditionals, 302, 317
- terms. *See* singular terms
- testimony, 218, 234, 239, 531–533, 704n22. *See also* assertion(s); deontic scorekeeping; inheritance of entitlements  
 authority of, 39–41, 175, 205, 234, 709n75
- 'that', 68, 70, 329, 501–503, 505–506, 535, 538–539, 598, 608, 705n30
- theoretical concepts. *See* concepts, theoretical

- theoretical reasoning, 116, 520–522. *See also* practical reasoning
- 'thing' as pseudosortal, 438
- thoughts, 154, 567–568, 570. *See also* assertion(s); intentional states; propositional contents; rationality; sapience; singular thoughts
- tokenings, 221, 303, 309–310, 432, 535, 537, 541, 566, 582, 664n10, 677n24, 685n40, 697n74, 705n26
- token-recurrence structures, 432, 451–458, 469, 470, 488–489, 503, 539, 566, 573, 698n80. *See also* substitution and anaphora, 432, 455, 456, 467, 472, 490, 499, 511, 580 asymmetric/symmetric, 455, 457, 490, 499, 564, 621–622, 706n33
- tonk, 125
- translation, 409–412, 535, 575
- transparency, 128, 292–293, 570, 576, 583
- triangulation, 426–432, 449, 467. *See also* reliable differential responsive dispositions
- 'true', xvii, 112, 116, 137, 202, 278, 280, 283–285, 288, 297–300, 302, 304–307, 321–322, 327, 336, 499, 516–517, 527–528, 530. *See also* truth expressive role of, 202, 278, 283–285, 324, 326–333, 568
- truth, xii, xvii, 5, 8, 11–13, 17, 80–82, 99, 111, 112, 202, 236, 240, 277, 288, 289, 291, 298–303, 317, 318, 322–324, 344, 346, 349, 490, 498, 515, 516, 526, 528, 568, 594–596, 646, 649. *See also* making-true; representation; semantics; taking-true conditions, 5, 6, 94, 129, 137, 278, 285, 326, 329, 356, 364, 513–517, 522, 599, 605, 669n89, 681n2, 710n95 correspondence theories of, 291, 317–318, 326, 330, 333 and *de re* ascriptions, 517, 595, 598, 646 explanatory role of, 286, 328–329 as expressive, 204, 232, 287, 299, 328–329, 333 and facts, 333, 622, 624, 625, 704n19 and inference, xvii, 5–6, 96–97, 104–105, 107–108, 277, 370, 519 performative account of, 287–288 as a property, 203, 303, 322, 323, 325–327 presentential theory of, 301–306, 322 and redundancy of force, 291, 299–300 and representation, 70–72, 568 and success, 236, 286, 287, 291, 521, 527–528 and taking-true, 287, 292, 298, 322, 599, 627 talk, 17–18, 278, 279, 322, 515 truth-values, 80–81, 340, 344, 345, 352, 355–357, 617, 669–670n1. *See also* designatedness trying, 294–295, 523–524. *See also* action(s) twin-earth, 119–120, 331–332, 703n18
- understanding, 4–5, 7, 13–14, 17, 31–32, 64–66, 74, 80, 90, 120, 215, 216, 294, 295, 478, 517, 526, 636. *See also* content(s); deontic scorekeeping; explicit; pragmatism and interpretation, 508, 510, 658n39, 660n56 and knowledge, 209, 213–215 and meaning, 60, 62, 73, 88–89, 478 undertaking, 167, 194, 598, 627. *See also* acknowledgment; commitment(s); deontic attitudes; deontic scorekeeping and acknowledgment, 194, 554 and attributing, 62, 161–163, 507, 554, 596–598, 677n11
- unexplained explainers, 133, 360
- unjustified justifiers, 221, 228. *See also* foundationalism
- unrepeatability, 282, 432, 454, 466, 495, 592. *See also* anaphora; token-recurrence structures
- validity, 10, 102–103, 117, 340, 345–346, 666n41, 687n8. *See also* formal inferences; logical vocabulary assertional/inferential, 347–348 definition of, 346, 383 and material inference, 131, 133–134, 347–348, 359 and substitution, 104–105, 135, 348, 383, 619
- verificationism, 121
- vindication, 174, 176–177, 193, 204, 218, 226, 239, 265, 532–533. *See also* deference; deontic scorekeeping; justification
- vocabulary. *See* logical vocabulary; normative vocabulary; semantic vocabulary

- volitions, 9, 258. *See also* practical commitments
- we, xi, xxi, 3–5, 7, 32, 39, 203, 275, 508, 587, 639, 641–645, 650, 716n36. *See also* demarcation; interpretation
- weak *de re* ascriptions. *See de re* ascriptions, weak
- Weber, Max, 48
- 'will', 245–246
- will, 32, 49, 233, 270, 271, 295, 661n69.  
*See also* action(s); akrasia; practical reasoning; trying
- Wittgenstein, Ludwig, 22, 52, 55, 74–76, 80–82, 111, 128, 178, 187, 226, 255–256, 333, 343, 461
- on interpretation, 61, 62, 508–510
- on norms, 13–16, 18, 46, 55, 200, 289, 603, 656n10
- pragmatism of, xii, 21–25, 29–30, 32, 42, 53, 55, 172, 199, 288
- on regress of rules, 20–22, 26, 45, 184, 206
- use of 'Regel', 64–66
- word-word relations, 306, 307, 325
- word-world relations, 306, 312–313, 324, 464
- Wright, Crispin, 37, 53–54, 656n10, 669n91