FREGE'S PUZZLE

Nathan Salmon

Frege's Puzzle
Nathan Salmon
Cambridge
The MIT Press
1986
xi, 194

Contents

Preface ix
Errata and Alterations xiii

Introduction 1

Chapter 1
Frege's Puzzle and the Naive Theory 11
1.1 *Frege's Puzzle and Information Content* 11
1.2 *The Naive Theory* 16

Chapter 2
Frege's Puzzle and the Modified Naive Theory 19
2.1 *The Singly Modified Naive Theory* 19
2.2 *The Doubly Modified Naive Theory* 24

Chapter 3
The Theories of Russell and Frege 45
3.1 *Russell* 45
3.2 *Frege* 46

Chapter 4
The Structure of Frege's Puzzle 55
4.1 *Compositionality* 55
4.2 *Frege's Law* 57
4.3 *Challenging Questions* 61

Chapter 5
A Budget of Nonsolutions to Frege's Puzzle 63
5.1 *Conceptual Theories* 63
5.2 *Contextual Theories* 70
5.3 *Verbal Theories* 71
5.4 *Frege's Strategy Generalized* 73

Chapter 6
The Crux of Frege's Puzzle 77
6.1 The Minor Premise 77
6.2 Substitutivity 79

Chapter 7
More Puzzles 87
7.1 The New Frege Puzzle 87
7.2 Elmer's Befuddlement 92

Chapter 8
Resolution of the Puzzles 103
8.1 Attitudes and Recognition Failure 103
8.2 Propositional Attitudes and Recognition Failure 105
8.3 Resolution 109
8.4 Why We Speak the Way We Do 114

Chapter 9
The Orthodox Theory versus the Modified Naive Theory 119
9.1 Semantics and Elmer's Befuddlement 119
9.2 Quantifying In 121
9.3 Propositional-Attitude Attributions 125
9.4 Concluding Remarks 126

Appendix A Kripke's Puzzle 129

Appendix B Analyticity and A Priority 133

Appendix C Propositional Semantics 143

Notes 153

Bibliography 181

Index of Theses 187

Subject Index 191

Preface

This book concerns topics that have occupied me since 1972, when I was an undergraduate. It began as a sketch for a paper in late 1980, when I was first struck by a cluster of arguments—some due to Keith Donnellan, some to David Kaplan, some to Saul Kripke, and some to me—that, taken collectively, finally convinced me of a philosophical thesis I had always thought to be fundamentally mistaken and to have been essentially refuted by Gottlob Frege: that the thoughts we have and the propositions we assert, believe, or bear other propositional attitudes toward, when formulatable using ordinary proper names, are always Russellian "singular propositions" (Kaplan), in which the only thing contributed by a name's occurrence is the named individual, and, furthermore, the attributions of thought, assertion, belief, and other attitudes we make using proper names do nothing more or less semantically (at the level of proposition content) than ascribe thought, assertion, belief, or other attitudes toward just such propositions. I was thus led to accept and defend the consequence that co-referential proper names are always intersubstitutable, salva veritate, in attitude contexts, as well as certain other unpopular consequences. There was to be solved one major and obvious philosophical difficulty with the thesis: It seemed decisively false. I soon discovered further consequences of and difficulties with the thesis. I also discovered that my idea for solving the major difficulty with it also yielded solutions, or partial solutions, to many of the other difficulties.

The work was expanded and revised between 1981 and 1985, in Princeton and California, amid a variety of circumstances that made sustained work impossible. Since 1980, when I was first struck by the arguments, I have remained firmly convinced of the thesis and its consequences.

Regrettably, it is beyond my powers of recall to thank everyone who has influenced my thinking on these topics. I would be remiss indeed, however, if I did not acknowledge my profound debt to my former teachers Tyler Burge, Alonzo Church, Keith Donnellan, David Kaplan,

and Saul Kripke. Each has influenced my thinking on these matters in a great many ways, though of course none need agree with all that I say here. Indeed, much of what I have to say is in sharp conflict with the views of Burge and Church, as I understand them. Kaplan has often expressed a strong inclination toward something like the thesis mentioned above, but he has also often expressed a reluctance to accept its more bitter consequences, and he has recently informed me that he disbelieves it. Donnellan and Kripke may be somewhat more sympathetic, but I believe that even they are somewhat uncomfortable with some of the views I defend here.

To Saul Kripke I owe an additional and special sort of debt, closely akin to my obvious debts to Gottlob Frege and Bertrand Russell. Through Kripke's penetrating and enlightening work—especially the marvelous paper "A Puzzle About Belief" and that newer paradigm of philosophy *Naming and Necessity*—I have come to see matters in a way that would scarcely be possible without them. In addition to being influenced by his informative and exciting published works and public lectures, my philosophical development has benefited immeasurably from my friendship with him. My views on the issues addressed in this book do not always coincide exactly with his; in fact, in the first two appendixes I criticize some of his published remarks on particular topics. But it remains that my general philosophical method and point of view owe a great deal to him and his work.

I am similarly indebted to David Kaplan, both for the many ideas that stem from his work and for the many intellectual benefits that accrue from being his friend.

The philosophical literature pertaining to my topic is immense and is expanding daily. No attempt is made here to discuss all or even a good portion of the significant works. Many of the most important contributions are cited only briefly in the notes, and some are not mentioned at all—especially some that have appeared in the years since 1980, during which my original sketch was expanded and revised. Some of the rest are discussed at slightly greater length, primarily in the notes. With some exceptions, the bibliography lists works that I actually consulted in writing the book; consequently, it does not constitute a complete list of important works on any of the subjects treated herein.

Portions of previous versions of this book were delivered as talks between 1982 and 1984 at a number of universities. The discussions that followed led to many improvements over former versions. I also benefited from the written comments and suggestions of David Austin, Hector-Neri Castañeda, Graeme Forbes, Ruth Barcan Marcus, Stephen Schwartz, and an anonymous reviewer, and from separate discussions

of particular aspects of the book with Joseph Almog, Blake Barley, John Biro, Francis Dauer, Keith Donnellan, Edmund Gettier, Gilbert Harman, Mark Johnston, David Kaplan, Saul Kripke, Igal Kvart, David Lewis, David Magnus, Ruth Barcan Marcus, Mark Richard, Howard Wettstein, and (especially) Scott Soames.

I am grateful also to Princeton University for allowing me a research leave in the spring of 1981, during which the first draft of what was to become the present book was written, to the philosophy departments of the University of California at Riverside and at Santa Barbara for generously allowing me the use of secretarial, word-processing, printing, and photocopying resources far in excess of my fair share, and to Angie Arballo at Riverside and Paula Ryan at Santa Barbara for their excellent typing. Finally, I am indebted to UCSB and Marilyn Freeman at the Raytheon Corporation in Santa Barbara for financial and technical assistance, respectively, in the preparation of the manuscript.

Errata and Alterations

Note: The hurried reader who is not interested in the details of the theory defended here (especially as they pertain to temporal issues concerning propositions), but only in how the theory is defended, should skip the long section 2.2, pp. 24-44.

('fb' abbreviates 'from bottom')

page	*line*	*Replace:*	*With:*
ix	4 fb	thank everyone	thank everyone by name
3	20 fb	intensional	nonextensional
4	10 fb	only semantic value	semantic content
7	19	compelling	forceful
14	1 fb	possble	possible
16	10 fb	the *'Fido-Fido'*	*the 'Fido-Fido'*
24	1 fb	it Let	it. Let
45	6 fb	more accurate	perhaps
46	3-4	uniquely instantiated	instantiated uniquely
79	14	Frege's law	Frege's Law
109	7	one the main	one of the main
115	14-12 fb	that the ancient astronomer believes (his sentence for) the sentence 'Hesperus is Phosphorus' to be true and, under normal circumstances, would verbally assent to it if queried.	that, assuming the ancient astronomer understands (his version of) the sentence 'Hesperus is Phosphorus', under normal circumstances he would verbally assent to it if queried.
129	3	propositional attitude	propositional-attitude

133	9 fb	sensory experience.	experience.
135	16 fb	Phosphorus,	'Phosphorus,
141	13 fb	tlw (content of	(the content of
145	9-10	*'Necessary)*	*'Necessary'*)
146	8, 10	*ι* [Greek iota]	ι [inverted Greek iota]
151	5 fb	*P*	*p*
161	8	are the object and the propositional function	include the object
163	3	circumstances	circumstance
168	20 fb	suspects that *S* theorizes	suspects that *S*, theorizes
170	4	to themodified	to the modified
172	1	the general idea	Kaplan's general idea
178	6	& *x* is a fox)	& *x* is a fox)'
178	22 fb	'fortnight may	'fortnight' may
179	13	... a fortnight ...	a fortnight is *n* days
182	13	647-658.	281-304.
182	15 fb	Gedankengefuge	Der Gedanke
183	3	Pierre Believes."	Pierre Does Not Believe."
184	15	(1979):	(1981):
184	14 fb	*Thoughts.*	*Thought.*
193	1 fb	Richard, M., 26, 158	Richard, M., 26, 157, 158

It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a terrestrial being for the very first time can be put into a form of words which will be understood by someone to whom the thought is entirely new. This would be impossible, were we not able to distinguish parts in the thought corresponding to the parts of a sentence, so that the structure of the sentence serves as an image of the structure of the thought. To be sure, we really talk figuratively when we transfer the relation of whole and part to thoughts; yet the analogy is so ready to hand and so generally appropriate that we are hardly ever bothered by the hitches which occur from time to time.

Gottlob Frege, opening paragraph of "Gedankengefüge"

Introduction

The topic of this book is the nature of the cognitive information content of declarative sentences such as 'Ted Kennedy is tall' and 'Saul Kripke wrote *Naming and Necessity*', as uttered in a particular possible context. My aim is to motivate and defend a certain sort of theory of content, one that has been rejected as patently false by the majority of contemporary philosophers of language. I shall argue that a certain version of the theory is true; however, given the controversial nature of the theory and given the nature of philosophy in general, the ultimate goal of this essay is simply to convince the reader that the theory I defend is at least as reasonable as any of its rivals.

The theory holds that the cognitive content of the sentence 'Ted Kennedy is tall', with respect to some context *c*, is a complex entity called a *proposition*, made up somehow of the man Kennedy, the attribute (property) of being tall, and the time of the context *c*, and that the content of the sentence 'Saul Kripke wrote *Naming and Necessity*', with respect to a context *c*, is made up of the man Kripke, the work *Naming and Necessity*, the attribute of authorship (i.e., the relation of having written), and the time of *c*. Propositions of this sort, in which individuals whom the proposition is about "occur as constituents" (to use Bertrand Russell's phrase), are what David Kaplan has called *singular propositions*. By contrast, a (*purely*) *general proposition* is a composite purely intensional entity made up solely of further intensional entities such as attributes and concepts, employing purely conceptual representations of the individuals whom the proposition is about in place of the individuals and times themselves. Such might be the content of a sentence like 'A certain sometimes popular legislator is often outspoken'. The great philosopher of mathematics and language Gottlob Frege maintained that the cognitive content (what he called the *Erkenntniswerte*) of any complete declarative sentence is always a purely general proposition, or what he called a 'thought' (*Gedanke*). I shall call any theory of this sort *Fregean*.

The theory of content that I shall defend is quite natural from a philosophical point of view and quite simple from a semantic point of

view. Yet it is, philosophically and semantically, quite powerful. Some of its consequences are surprising and may seem counterintuitive—for example, its consequences concerning analyticity and syntheticity, a priority and a posteriority. On the traditional view, sentences like 'The planet Neptune, if it exists, causes perturbations in the orbit of Uranus' and 'The Standard Meter, if it exists, is exactly one meter long' are analytic, and their content is therefore both necessary and knowable solely by a priori means, whereas sentences like 'Hesperus, if it exists, is Phosphorus' and 'If Cicero was a Roman orator, then so was Tully' are synthetic, contingent, and a posteriori. In 1970 Saul Kripke astonished the analytic philosophical community with his claim—supported by the rich theoretical apparatus of possible-world semantics and his new "picture" of reference—that the former two sentences, though a priori, are in fact contingent and therefore synthetic, whereas the latter two sentences, though synthetic and a posteriori, contain necessary truths, propositions true in every possible world. The theory I advance here accords entirely with Kripke's view that the former two sentences are contingent but contends that the second at least is also a posteriori. Furthermore, it agrees that the latter two sentences are necessary but contends that they are also a priori. In fact, I shall hold that they are analytic. This aspect of the theory is developed in an appendix. Part of my aim in this book is to make these and other surprising consequences palatable, by showing that the theoretical postulates that generate them are, in fact, in perfect accord with our intuitions concerning the cognitive information content of declarative sentences.

It is sometimes argued, and more often taken for granted, that the theory of singular propositions is, from the point of view of cognitive psychology, wholly inadequate and wildly implausible as a theory of the content of thought. The main idea behind this objection to the sort of theory I advocate might be illustrated by the following sort of thought experiment: Suppose Tom, Dick, and Harry, who have never met one another, agree to think some simple thought. Their instructions are 'Think to yourself that Ted Kennedy is tall', and each complies. Surely what goes on in each thinker's mind will differ considerably from one thinker to the next, varying with the thinker's political ideology and his familiarity with Kennedy's physical appearance, achievements, deeds, and so on. Tom thinks something along the lines of "That famous senator from Massachusetts is tall", while Dick thinks "That handsome brother of Jack and Bobby is tall", while Harry thinks "That good-for-nothing !@%!@ is tall". As each apprehends the words 'Ted Kennedy is tall', the content of his thought is something much richer, in structure and thought-stuff, than the crude singular proposition postulated by the theory in question. Each thinker thinks a different

thought. Of course, these various thoughts, though different in content, are not completely and utterly dissimilar; otherwise the thinkers in our experiment could hardly be said to be unanimously complying with their instruction to think that Ted Kennedy is tall. Though Tom, Dick, and Harry are thinking different thoughts, they have in common that each thinks a thought *about* the man Kennedy, to the effect that he is tall. To use the familiar locution of so-called *de re thought*, though they are thinking different thoughts,

(1) Tom is thinking of Ted Kennedy that he is tall,

and so is Dick and so is Harry. The more formally inclined would put it thus:

(2) $(\exists x)[x =$ Ted Kennedy & Tom is thinking that x is tall],

and similarly for Dick and Harry. To assert these things, the objection continues, is not to say that there is a special, crudely individuated thought content that each thinker shares. De re locutions such as those in examples 1 and 2 do not *specify* fully a particular content apprehended by the subject, nor do they pretend to. They merely *characterize* a content, by specifying what kind of content it is. The critical feature of attributions like 1 and 2—what makes them de re rather than de dicto—is that the proper name 'Ted Kennedy' is positioned outside the oblique context created by the intensional operator 'Tom is thinking that', where it is open to substitution of co-referential singular terms and to existential generalization. In Quine's terminology, what makes the attributions 1 and 2 *relational* rather than *notional* is the fact that the name 'Ted Kennedy' occurs within a *transparent context*, outside the *opaque* context, and is therefore in *purely referential position*. In order to attribute a particular content to Tom—in order to specify his thought content and not merely characterize it—one would need a suitable singular term occurring within the oblique context (in Quinean theory, the opaque context) created by 'Tom is thinking that', as in

Tom is thinking that: the senior senator from Massachusetts is tall.

To suppose that attributions 1 and 2 actually specify a particular content, the objection continues, is to misread a de re attribution of thought as if it were something else altogether: an attribution of a peculiar and special sort of thought, a brand new kind of beast. Thus, the psychologically unpalatable postulation of singular propositions as the content of thought is seen as a fallacious inference based on a misunderstanding of de re attributions such as 1 and 2.[1]

Singular propositions as the contents of psychological states are not so easily dismissed. This common objection to the theory of singular

propositions fails to appreciate some of the finer points of contemporary formal semantic analysis, finer points which lead directly in this case to philosophical illumination.[2] It was argued that the critical feature of de re attributions such as 1 and 2 is that the name 'Ted Kennedy' is positioned outside the scope of 'Tom is thinking that'. What is more significant, however, is that some other singular term occurs *within* the nonextensional (or "oblique" or "opaque") context—the anaphoric pronoun 'he' in example 1, the variable '*x*' in example 2. Consider first the quasi-formal sentence 2. Once it is granted that this sentence is true, it follows by principles of conventional formal semantics that its component open sentence

(3) Tom is thinking that *x* is tall

must be true under the assigment of Ted Kennedy as the value of the variable '*x*'—in the terminology of Tarski, that Ted Kennedy *satisfies* sentence 3. Similarly, on the less formal rendering 1, its component sentence

(4) Tom is thinking that he is tall

must be true under the anaphoric assignment of Kennedy as referent for the pronoun 'he'. The open sentence 3 is true under the assignment of Kennedy as the value of '*x*' only if Tom is thinking (i.e., having a thought whose content is) the semantic content of the embedded open sentence

x is tall

under the same assignment of Kennedy as the value of '*x*'. Similarly, sentence 4 is true under the assignment of Kennedy as referent for the pronoun 'he' only if Tom is thinking the semantic content of

he is tall

under this same assignment. Now, the fundamental semantic characteristic of a variable '*x*' with an assigned value, or of a pronoun 'he' with a particular referent, is that its only semantic value is its referent. There is nothing else for it to contribute to the semantic content of the sentences in which it figures. Indeed, this is precisely the point of using a variable or a pronoun within the scope of the attitude verb in a de re attribution. If the variable or pronoun had, in addition to its referent, something like a Fregean *sense*—something conceptual that it contributed to semantic content—the speaker's intention of declining to specify the way in which the subject of the attribution thinks or conceives of the res in question would be thwarted, and the attribution would be de dicto instead of de re, notional instead of relational. Thus, the content

of '*x* is tall', or 'he is tall', under the assignment of Kennedy as referent for '*x*' or 'he', can only be the singular proposition about Kennedy that he is tall, the crudely individuated sort of entity postulated by the theory I am advocating. And each thinker thinking de re of Kennedy that he is tall thinks this same content.

The de re locutions 1 and 2 do more than merely characterize the content of the attributed thought; they fully and uniquely specify a particular content after all. More important, even if one allows only that there is someone—not specified—such that Tom is thinking that he is tall, or, more formal,

(∃*x*) Tom is thinking that *x* is tall,

it still follows by principles of formal semantics that Tom is thinking the semantic content of '*x* is tall' under some assignment or other to '*x*'—or the semantic content of 'he is tall' under some assignment or other to 'he'—and hence that Tom is having a thought whose content is some unspecified singular proposition. Once it is granted by way of the familiar de re locution that someone is apprehending the thought (or believes, or knows) of someone or something that he, she, or it has a certain property, it follows by the principles of formal semantics that the subject of the de re attribution apprehends (believes, knows) a singular proposition. The move to de re locutions in characterizing what various thinkers have in common does not refute the theory of singular propositions as contents of thought; on the contrary, it proves the theory.

If there is any lacuna in this semantic argument for singular propositions, it is in the inference from the observation that the open sentence 3 is true under a particular assignment of a value to the variable '*x*' to the conclusion that Tom is thinking the semantic content of '*x* is tall' under that same assignment. This inference is validated by a simple and straightforward analysis of the sentential form ⌜*a Vs* that *S*⌝, where *V* is any of a certain class of attitude verbs, including 'think', 'believe', 'know', 'say', and many more. On the most plausible construal, a sentence of this form is an atomic predication, with *a* a singular term, ⌜*Vs*⌝ a dyadic predicate, and ⌜that *S*⌝ a second term much like a singular term. (This construal best explains the inference from 'Jones believes that all men are created equal' to 'There is something that Jones believes but Smith doubts', given the further premise 'Smith doubts that all men are created equal', or to 'Jones believes the proposition to which our nation is dedicated', given the further premise 'That all men are created equal is the proposition to which our nation is dedicated'.) Now it is necessary to inquire into the nature of a propositional term of the form ⌜that *S*⌝, with *S* a sentence. The word 'that', in its use as

forming 'that'-clauses from sentences, is best regarded as a propositional-term-forming sentential operator, not unlike a pair of quotation marks. Of course, semantically the 'that'-operator differs significantly from quotation marks. The semantics of quotation marks is more or less exhausted by the following rule:

> The result of enclosing any expression within quotation marks refers, with respect to semantic parameters (such as a time, a possible world, a context, or an assignment of values to variables) to the enclosed expression itself.

The semantics of the 'that'-operator is governed by a somewhat different rule:

> For any (open or closed) sentence S, the result of prefixing S with the 'that'-operator, that $\bar{S}$, refers with respect to semantic parameters (such as a time, a possible world, a context, or an assignment of values to variables) to the semantic content of S with respect to those parameters.

Given this semantic rule, the inference in question goes through. Hence, so does the semantic argument for singular propositions.

The semantic argument shows that singular propositions are the contents of thoughts and beliefs, that we have propositional attitudes toward singular propositions. But the argument seems to show much more than this. It appears to yield the result that anyone who is thinking (anyone who believes, etc.) that Ted Kennedy is tall thinks (believes, etc.) the singular proposition about Kennedy that he is tall. In light of the semantic rule for 'that'-clauses given above, this result may be regarded as confirmation of the sort of theory of semantic content that I am advocating, the theory that the semantic content of the sentence 'Ted Kennedy is tall' is just the singular proposition. The semantic argument thus appears to refute the Fregean theory. At the very least, it deals a heavy blow to the spirit of that theory.[3]

The argument does not literally disprove the Fregean theory. An alternative interpretation is available—one that is consistent with the letter, if not the spirit, of the Fregean theory, and independently plausible in its own right. One's thought is *of* a particular object, it may be argued, only by virtue of one's grasping some general proposition related to the object in a special way. Even if it must be allowed that we have thoughts and beliefs whose contents are singular propositions, to have such a thought or belief just *is* to have a thought or belief whose content is a general proposition, one that determines, in a certain way, the singular proposition. Tom thinks of Kennedy that he is tall by thinking a general proposition, Dick by thinking another general proposition,

and Harry by thinking yet another. The sentence 'he is tall' under the assignment of Kennedy as referent for the pronoun 'he' expresses a singular proposition, it will be granted; however, as is evident from the formalized rendering 2 of the de re attribution 1, this is because the sentence 'he is tall', occurring as the 'that'-clause in a de re attribution such as 1, functions like an open sentence of formal logic: '*x* is tall' with free '*x*'. The sentence 'Ted Kennedy is tall', on the other hand, is a closed sentence; it contains no free variable with an assigned value, no pronoun with anaphoric reference, only the proper name 'Ted Kennedy' (a closed singular term). The semantic argument shows that anyone who is thinking (believes) of Kennedy that he is tall thinks (believes) the singular proposition about Kennedy that he is tall. It does not prove that this same singular proposition is the cognitive information content of the closed sentence 'Ted Kennedy is tall'.

It may seem a hollow triumph for the Fregean theory if singular propositions are eschewed as the cognitive contents of closed sentences involving proper names only to be foisted on us as the cognitive contents of the thoughts we have when we understand these sentences. There are, however, compelling Fregean reasons for rejecting singular propositions as the contents of closed sentences. Some of these reasons concern problems that arise in connection with sentences involving names that do not refer to anything, such as perhaps those found in works of fiction. I shall mention the problems concerning nonreferring singular terms only briefly in this book. The general problem of nonreferring singular terms is large and complex and could well serve as the topic of another book. In this book, however, I shall be primarily concerned with another nest of problems stemming from a puzzle due to Frege concerning the cognitive content of statements, especially identity statements. This has been variously called 'the problem of the Morning Star and the Evening Star', 'the Hesperus-Phosphorus problem', and 'Frege's puzzle about identity', among other things. I shall call it simply 'Frege's Puzzle'. I shall propose a reformulation of the puzzle and sketch a solution. In doing so, I shall discuss various aspects of the puzzle: what it is a puzzle about; what makes it a puzzle; what, if anything, the puzzle shows; why various proposals for solving the puzzle, both new and old, do not succeed; and what a satisfactory solution to the puzzle ought to be like. My proposal does not solve all the philosophical questions that arise in connection with the general problem, but it is my hope that my efforts to dilute the general problem will strengthen the prospects for solutions to some of the remaining problems.

The theory of singular propositions as the contents of closed declarative sentences is more or less explicit in the writings of Bertrand

Russell. The problems stemming from phenomena such as Frege's Puzzle and nonreferring singular terms led Russell to retreat from the simple-minded version of the theory, according to which a closed sentence involving a proper name of an ordinary individual, e.g. 'Ted Kennedy is tall', conveys a singular proposition about that individual as its cognitive content. Denying that it is possible to grasp a singular proposition involving an ordinary, external object as a constituent, Russell held that in order to understand such a sentence one is compelled to construe it as containing a different proposition, one not involving the relevant object directly as a constituent. My purpose is to develop and defend the simple-minded version of the theory that Russell came to reject. I shall suggest a way of extending the theory in such a way as to deal with the difficulties generated by Frege's Puzzle and other problems of its ilk. In this respect, the version of the theory I shall propose is similar in outline to certain theories advanced by others recently working in the same area, most notably David Kaplan (in "Demonstratives," draft no. 2) and John Perry (in "The Problem of the Essential Indexical"). There are significant differences in detail, however, and I arrive at my destination via a route somewhat different from those taken by these philosophers. In fact, both Kaplan and Perry have recently expressed views that conflict sharply with some aspects of the view I shall advocate. My purpose here is not to invent an entirely new and original theory of reference and cognitive information content, but to develop arguments for a certain theory that is familiar in broad outline, to develop this theory in some detail, to uncover some of its important but generally unnoticed consequences, and to make these consequences palatable.

In addition to Frege's Puzzle I shall discuss the related problem of the apparent failure of substitutivity of co-referential proper names in propositional attitude attributions. What I have to say here has straightforward implications for related problems, including Kripke's puzzle about belief, the problem of nested propositional attitude contexts posed by Benson Mates, and the problems of thought about oneself posed by Hector-Neri Castañeda. Kripke's puzzle is discussed in an appendix; the problems posed by Castañeda and Mates are discussed, sketchily, in notes to relevant passages.

There is a third appendix, in which an extremely primitive and elementary semantics of singular propositions is outlined in accordance with the theory defended here. The purpose is to give the reader some rudimentary idea of one direction, among many, that a detailed formal development might take. Most of the interesting technical questions that naturally arise are not addressed here, but it is hoped that the

primitive outline conveys some idea of what complex propositions might "look like," how they may be true or false, how a set of them may collectively entail another, how they may be individually valid or invalid, and so on.

Chapter 1
Frege's Puzzle and the Naive Theory

1.1 *Frege's Puzzle and Information Content*

> Identity challenges reflection through questions which are connected with it and are not altogether easy to answer. . . . '$a = a$' and '$a = b$' are obviously sentences of a different cognitive value [Erkenntniswerte]: '$a = a$' holds a priori and is according to Kant to be called analytic, whereas sentences of the form '$a = b$' often contain very valuable extensions of our knowledge and are not always to be grounded a priori. . . . If we then wanted to view identity as a relation between that which the names *a* and *b* signify, then '$a = b$' and '$a = a$' would seem to be potentially not different, in case, that is, '$a = b$' is true. There would be thereby expressed a relation of a thing to itself, one in which each thing stands to itself, but no thing stands to another.

With the German equivalent of these words, Frege tortuously poses the problem that gave rise to his celebrated theory of sense: How can '$a = b$', if true, differ in "cognitive value"—that is, in cognitive information content—from '$a = a$'? Clearly they differ, since the first is informative and a posteriori where the latter is uninformative and a priori. But, assuming that '$a = b$' predicates the relation of identity between the referent of the name *a* and the referent of the name *b*, and that '$a = a$' predicates the relation of identity between the referent of *a* and the referent of *a*, then if '$a = b$' is true, it predicates the same relation between the same pair of objects as does '$a = a$'. It would seem, then, that '$a = b$' and '$a = a$' ought to convey the same piece of information. But clearly they do not. So what gives here?

A number of philosophers have found the identity relation, taken as the relation that "each thing stands in to itself, but no thing stands to another," curious, mysterious, or bogus. In the *Tractatus* (sections 5.53–5.535), Wittgenstein denies that there is any such relation.[1] Earlier, in *Begriffsschrift* (section 8), Frege took a similar tack, proposing an analysis of identity sentences according to which singular terms "display

their own selves [appear in propria persona] when they are combined by means of the sign ['='] for identity of content [referent], for this expresses the circumstance of two names [singular terms] having the same content [referent]." Thus the early Frege and Wittgenstein attempted to rid themselves of the puzzle. More recent philosophy has followed Frege's later characterization of the origins of the puzzle as one arising from reflection on the concept of the identity by the use of such epithets as 'Frege's puzzle about identity' or 'Frege's identity problem'. The first point I wish to emphasize about 'Frege's puzzle about identity' is that, pace Frege, it is not a puzzle about identity. It has virtually nothing to do with identity. Different versions of the very same puzzle, or formally analogous puzzles that pose the very same set of questions and philosophical issues in the very same way, arise with certain constructions not involving the identity predicate or the identity relation. For example, the sentence 'Shakespeare wrote *Timon of Athens*' is informative, whereas 'The author of *Timon of Athens* wrote *Timon of Athens*' is not. The same question arises: How can that be? Given that the first sentence is true, it would seem that both sentences contain the same piece of information; they both attribute the same property (authorship of *Timon of Athens*) to the same individual (Shakespeare). This kind of example is unlike Frege's version of the puzzle in that it involves a definite description, whereas Frege's can involve two proper names and consequently applies pressure against a wider range of semantic theories. It is not difficult, however, to construct further puzzling examples involving two names without using the identity predicate; the sentence 'Hesperus is a planet if Phosphorus is' is informative and apparently a posteriori, whereas the sentence 'Phosphorus is a planet if Phosphorus is' is uninformative and a priori. However, both sentences attribute the same property, *being a planet if Phosphorus is*, to the same entity, the planet Venus. Looked at another way, both sentences attribute the same relation, *x is a planet if y is*, to the same (reflexive) pair of objects. In either case, the two sentences seem to contain the very same information.

It is easy to see from these examples that versions of Frege's Puzzle can be constructed in connection with any predicate whatsoever, not just with the identity predicate. What, then, is the general puzzle about if it is not a puzzle about identity? These same examples provide the answer. The general problem is a problem concerning pieces of information (in a nontechnical sense), such as the information that Socrates is wise or the information that Socrates is wise if Plato is. The various versions of Frege's Puzzle are stated in terms of declarative sentences rather than in terms of information. This is because there is an obvious and intimate relation between pieces of information (such as the in-

formation that Socrates is wise) and declarative sentences (such as 'Socrates is wise'). Declarative sentences have various semantic attributes: they are true, or false, or neither; they have semantic intentions (i.e., correlated functions from possible worlds to truth values); they involve reference to individuals, such as Socrates; and so on. But the fundamental semantic role of a declarative sentence is to encode information.[2] I mean the term 'information' in a broad sense to include misinformation (that is, inaccurate or incorrect pieces of information), and even pieces of information that are neither true nor false. Pragmatically, we use declarative sentences to communicate or convey information to others (generally, not just the information encoded by the sentence), but we may also use declarative sentences simply to record information for possible future use, and perhaps even to record information with no anticipation of any future use. If for some reason I need to make a record of the date of my marriage, say to recall that piece of information on a later occasion, I can simply write the words 'I was married on August 28, 1980', or memorize them, or repeat them to myself. Declarative sentences are primarily a means of encoding information, and they are a remarkably efficient means at that. Many of their other semantic and pragmatic functions follow from or depend upon their fundamental semantic role of encoding information.

This statement of the semantic relation between declarative sentences and information is somewhat vague, but it is clear enough to convey one of the fundamental presuppositions of Frege's Puzzle. Vague though it may be, it is also obviously correct. Any reasonable semantic theory for declarative sentences ought to allow for some account of declarative sentences as information encoders, at least to the extent of not contradicting it. A conception of sentences as information encoders will be assumed throughout this book. A declarative sentence will be said to *contain* the information it encodes, and that piece of information will be described as the *information content* of the sentence.

Pieces of information are, like the sentences that encode them, abstract entities. Many of their properties can be "read off" from the encoding sentences. Thus, for instance, it is evident that pieces of information are not ontologically simple, but complex. The information that Socrates is wise and the information that Socrates is snub-nosed are both, in the same way, pieces of information directly about Socrates; hence, they must have some component in common. Likewise, the information that Socrates is wise has some component in common with the information that Plato is wise, and that component is different from what it has in common with the information that Socrates is snub-nosed. Correspondingly, the declarative sentence 'Socrates is wise' shares certain syntactic components with the sentences 'Socrates is snub-nosed'

and 'Plato is wise'. These syntactic components—the name 'Socrates' and the predicate 'is wise'—are separately semantically correlated with the corresponding component of the piece of information encoded by the sentence. Let us call the information component semantically correlated with an expression the *information value* of the expression. The information value of the name 'Socrates' is that which the name contributes to the information encoded by such sentences as 'Socrates is wise' and 'Socrates is snub-nosed'; similarly, the information value of the predicate 'is wise' is that entity which the predicate contributes to the information encoded by such sentences as 'Socrates is wise' and 'Plato is wise'. As a limiting case, the information value of a declarative sentence is the piece of information it encodes, its information content.

Within the framework of so-called possible-world semantics, the information value of an expression determines the semantic *intension* of the expression. The intension of a singular term, sentence, or predicate is a function that assigns to any possible world w the *extension* the singular term, sentence, or predicate takes on with respect to w. The extension of a singular term (with respect to a possible world w) is simply its *referent* (with respect to w), i.e., the object or individual to which the term refers (with respect to w). The extension of a sentence (with respect to w) is its truth value (with respect to w)—either truth or falsehood. The extension of an n-place predicate (with respect to w) is the predicate's semantic characteristic function (with respect to w), i.e., the function that assigns either truth or falsehood to an n-tuple of individuals, according as the predicate or its negation applies (with respect to w) to the n-tuple. Assuming bivalence, the extension of an n-place predicate may be identified instead with the class of n-tuples to which the predicate applies.

Since ordinary language includes so-called indexical expressions (such as 'I', 'you', 'here', 'now', 'today', 'yesterday', 'this', 'that', 'he', 'she', 'there', and 'then'), the information value of an expression, and hence also the semantic intension, must in general be indexed, i.e., relativized, to the context in which the expression is uttered. That is, strictly one should speak of the information value of an expression (e.g. the information content of a sentence) with respect to this or that context of utterance, and similarly for the corresponding semantic intension of an expression; the information value and corresponding intension of an expression with respect to one context may be different from the information value and corresponding intension of the same expression with respect to a different context. This generates a higher-level, non-relativized semantic value for expressions, which David Kaplan calls the *character* of an expression. The character of an expression is a function or rule that determines, for any possble context of utterance

c, the information value the expression takes on with respect to c. For example, the character of a sentence is a function or rule that assigns to any possible context of utterance c the piece of information that the sentence encodes with respect to c, that is, the information content of the sentence with respect to c.

In addition to the character of an expression, we may consider a related nonrelativized semantic value: the function or rule that determines for any possible context of utterance c the extension (e.g. the referent, the class of application, or the truth value) that the expression takes on with respect to c. Let us call this the *contour* of an expression. The contour of an expression is fully determined by its character, as follows: Given any context c, the character of an expression determines the information value of the expression with respect to c. This, in turn, determines the intension of the expression with respect to c. Applying this intension to the possible world of the context c yields the extension of the expression with respect to c.[3]

In summary, the central and fundamental semantic value of a declarative sentence is its information content, the piece of information encoded. This generates a fundamental semantic value for expressions generally: information value. The information value of an expression determines the expression's semantic intension, which assigns to any possible world a lower-level semantic value for the expression, its extension with respect to that possible world. Since ordinary language includes indexicals, the information value of an expression must be indexed to a context of utterance. This generates a higher-level semantic value for expressions, character, which assigns to any possible context the information value the expression takes on with respect to that context. The character of an expression determines the expression's contour, which assigns to any possible context the extension the expression takes on with respect to that context.

The systematic method by which it is secured which information is semantically encoded by which sentence is, roughly, that a sentence semantically encodes that piece of information whose components are the information values of the sentence parts, with these information values combined as the sentence parts are themselves combined to form the sentence.[4] In order to analyze the information encoded by a sentence into its components, one simply decomposes the sentence into its information-valued parts, and the information values thereof are the components of the encoded information. In this way, declarative sentences not only encode but also codify information.

One may take it as a sort of general rule or principle that the information value of any compound expression, with respect to a given context of utterance, is made up of the information values, with respect

to the given context, of the information-valued components of the compound. This general rule is subject to certain important qualifications, however, and must be construed more as a general guide or rule of thumb. Exceptions arise in connection with quotation marks and similar devices. The numeral '9' is, in an ordinary sense, a component part of the sentence 'The numeral '9' is a singular term', though the information value of the former is no part of the information content of the latter. I shall argue below that, in addition to quotation marks, there is another important though often neglected class of operators that yield exceptions to the general rule in something like the way that quotation marks do.[5] Still, it may be correctly said of any English sentence free of operators other than truth-functional connectives (e.g. 'If Socrates is wise, then so is Plato') that its information content is a complex made up of the information values of its information-valued components.

It is out of the natural and preliminary analysis presented here of the information contained in (i.e., semantically encoded by) declarative sentences, and not from reflection on the alleged mystique of identity, that Frege's challenging question arises.

1.2. *The Naive Theory*

What makes Frege's challenging question a puzzle? It poses a serious difficulty for a certain type of semantic theory: one that entails that two sentences involving an n-place predicate and n singular terms have the same information content (encode the same piece of information) if their predicates are semantically correlated with the same attribute (property or relation) and their singular terms, taken in sequence, coincide in reference, respectively. Specifically, the question poses a serious problem for what I, following David Kaplan, shall call the *naive theory*.

The naive theory is a theory of the information values of certain expressions. According to the naive theory, the information value of a singular term, as used in a possible context, is simply its referent in that context. This is similar to what Gilbert Ryle called the *'Fido'-Fido theory*, according to which the "meaning" or content of a singular term is simply its referent. Elements of this theory can be traced to ancient times. Likewise, the information value of a predicate, as used in a particular context, is identified with something like the semantically associated attribute with respect to that context, that is, with the corresponding property in the case of a monadic predicate or the corresponding n-ary relation in the case of an n-place predicate. On the naive theory, an atomic sentence consisting of an n-place predicate Π and n occurrences of singular terms, $a_1, a_2, \ldots, a_n$, when evaluated

with respect to a particular possible context, has as its cognitive content in that context a piece of information, called a *proposition*, which is supposed to be a complex consisting of something like the attribute referred to by Π with respect to that context and the sequence of objects referred to by the singular terms with respect to that context. For example, the cognitive information content of the sentence 'Socrates is wise' is to be the singular proposition consisting of Socrates and wisdom. On the naive theory, a sentence is a means for referring to its information content by specifying the components that make it up. A sentential connective may be construed on the model of a predicate. The information value of a connective would thus be an attribute (a property if monadic, a relation if polyadic)—not an attribute of individuals like Socrates, but an attribute of pieces of information, or propositions. For example, the information value of the connective 'if and only if' might be identified with the binary equivalence relation between propositions having the same truth value. Similarly, the information value of a quantifier might be identified with a property of properties of individuals. For example, the information value of the unrestricted universal quantifier 'everything' may be the (second-order) property of being a universal (first-order) property, i.e., the property of being a property possessed by every individual. The information value of a sentence, as used in a particular context, is simply its information content, the proposition made up of the information values of the information-valued sentence components.

The naive theory of information value, then, may be thought of as flowing from the following theses:

Thesis I
(Declarative) sentences encode pieces of information, called *propositions*. The proposition encoded by a sentence, with respect to a given context, is its *information content* with respect to that context.

Thesis II
The information content, with respect to a given context, of a sentence is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the sentence, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices.

Thesis III
The information value (contribution to information content), with respect to a given context c, of any singular term is its referent with respect to c (and the time of c and the world of c).

Thesis IV
Any expression may be thought of as *referring*, with respect to a given context, time, and possible world, to its information value with respect to that context.

Thesis V
The information value, with respect to a given context, of an *n*-place first-order predicate is an *n*-place attribute (either a property or an *n*-ary relation)—ordinarily an attribute ascribed to the referents of the attached singular terms. Exceptions arise in connection with quotation marks and similar devices.

Thesis VI
The information value, with respect to a given context, of an *n*-adic sentential connective is an attribute, ordinarily of the sorts of things that serve as referents for the operand sentences.

Thesis VII
The information value, with respect to a given context, of an *n*-adic quantifier or second-order predicate is an *n*-ary attribute, ordinarily of the sorts of things that serve as the referents for the operand first-order predicates.

Thesis VIII
The information value, with respect to a given context, of an operator other than a predicate, a connective, or a quantifier is an appropriate attribute (for sentence-forming operators) of, or operation (for other types of operators) on, the sorts of things that serve as referents for its appropriate operands.

Thesis IX
The information value, with respect to a given context, of a sentence is its information content, the encoded proposition.

Within the framework of the naive theory, the *meaning* of an expression might be identified with the expression's character, i.e., the semantically correlated function from possible contexts of utterance to information values. For example, the meaning of the sentence 'I am busy' will be thought of as a function that assigns to any context of utterance c the singular proposition composed of the agent of the context c (= the referent of 'I' with respect to c) and the property of being busy.

Chapter 2

Frege's Puzzle and the Modified Naive Theory

2.1 *The Singly Modified Naive Theory*

The naive theory is, as its name suggests, a prototheory of information value. For all its naiveté, there is a great deal to be said in its favor. First and foremost, it is a natural and compelling result, perhaps *the* natural result, of a preliminary philosophical investigation into the nature and structure of information. Some of the great thinkers in the philosophy of language, among them Frege and Russell, came to the subject with an initial presupposition of some rudimentary form of the naive theory. The theory yields a plausible rendering of the claim that the proposition that Socrates is wise is information *about* or *concerning* Socrates: The proposition is about Socrates in the straightforward sense that Socrates is an individual constituent of it. The naive theory extends easily to more complex sentential structures involving variables, connectives, quantifiers, and propositional operators. It gives substance to the oft-repeated slogan that to give (or to know) the semantic content (or "meaning," in the sense of information content) of a sentence or statement is to give (know) its truth conditions. Its notion of information content is exemplary of the kind of notion of proposition that is needed in connection with questions of de re modality: If I utter the sentence 'Socartes is wise', I assert something that is true if and only if the individual Socrates has the property wisdom. Moreover, what I assert is such that it is true with respect to an arbitrary possible world w if and only if that same condition, the very individual Socrates having wisdom, obtains in w. It is not enough, for instance, that someone in w who resembles or who represents the actual Socrates in a certain way be wise in w, or that someone in w who fits a certain conceptual representation of the actual Socrates be wise in w. It must be Socrates, the very individual. The naive theory also yields a straightforward notion of de re belief, and other de re propositional attitudes: To believe that p is to believe the proposition that p. So to believe *of* or *about* Socrates that *he* is wise is to believe the proposition *of* or *about* Socrates that *he* is wise, that is, the piece of information consisting of Socrates

and his wisdom. Indeed, as I argued in the introduction, these considerations concerning de re modality and de re propositional attitudes constitute important considerations favoring the naive theory over its rival, the orthodox Fregean theory, as well as over the theory of Russell.

Perhaps the most important thing to be said for the naive theory is that it has cogency and intuitive appeal as a theory of assertion. When I utter 'Socrates is wise', my speech act divides into two parts: I single someone out (Socrates), and I ascribe something to him (wisdom). These two component speech acts, singular reference and ascription, correspond to two components of *what I assert* when I assert that Socrates is wise. My asserting that Socrates is wise consists in my referring to Socrates and my ascribing wisdom to him; so too, that Socrates is wise (what I assert) consists of Socrates (what I refer to) and wisdom (what I ascribe to him).

So compelling is the naive theory that even Fregeans sometimes unconsciously and implicitly assume something like it. As I shall try to show, this may even be true of Frege himself when he argues against the naive theory in the very first paragraph of "Über Sinn und Bedeutung."

Compelling though it is, the naive theory has two fundamental flaws and must be modified if it is to yield a viable theory of information value. The first flaw is that the naive theory is in a certain sense internally inconsistent; the second concerns the eternalness of information. I shall consider each of these problems in turn.

The naive theory rests upon two central ideas. The first is the identification of the information value of a singular term with its referent, i.e., the 'Fido'-Fido theory (thesis III). By analogy, the referent of a predicate, a connective, or a quantifier is identified with its information value: the semantically correlated attribute of individuals, of propositions, or of properties of individuals, respectively (theses IV–VII). The second major idea is that the information value of a sentence, as uttered on a particular occasion, is made up of the information values of its information-valued components (theses II and IX). Unfortunately, these two ideas come into conflict in the case of definite descriptions. According to the naive theory, the information value of a definite description such as 'the individual who wrote *The Republic*' is simply its referent, Plato. Consequently, the sentence 'The individual who wrote *The Republic* is wise' is alleged to encode the singular proposition about Plato that he is wise. But the definite description is a phrase that, like a sentence, has parts with identifiable information values—for example, the dyadic predicate 'wrote' and the singular term (book title) '*The Republic*', as well as the monadic predicate 'wrote *The Republic*'. These information-valued components of the definite description are, ipso

facto, information-valued components of the containing sentence. If the information value (= information content) of a sentence is made up of the information values of its information-valued components, the information values of these description components must also go in to make up part of the information that the author of *The Republic* is wise. And if the information value of a sentence is something made up of the information values of its information-valued components, it stands to reason that the information value of a definite description, which is like a sentence at least in having information-valued components, should also be something made up of the information values of those components. Thus, instead of identifying the information value of 'the individual who wrote *The Republic*', as used on a particular occasion, with its referent, one should look instead for some complex entity made up partly of the relational property of having written *The Republic* (which, in turn, is made up of the binary relation *having written* and the work *The Republic*) and partly of something else—something that serves as the information value of the definite description operator 'the'. On this modification of the naive theory, the information that the author of *The Republic* is wise is not the singular proposition about Plato that he is wise but a different piece of information, one that does not have Plato as a component and has in his place something involving the property of authorship of *The Republic*.

Let us call this corrected version of the original theory the *singly modified naive theory*. One extremely important wrinkle in the singly modified naive theory is that a definite description 'the ϕ', in contrast with other sorts of singular terms, is seen as involving a bifurcation of semantic values taken on with respect to a context of utterance. On the one hand there is the description's referent, which is the individual to which the description's constitutive monadic predicate (or open formula) ϕ applies if there is only one such individual and is nothing otherwise. On the other hand there is the description's information value, which is a complex made up, in part, of the information value of the predicate (or formula) ϕ. By contrast, a proper name or other single-word singular term is seen as involving a collapse of semantic values; its information value with respect to a particular context is just its referent with respect to that context. From the point of view of the singly modified naive theory, the original naive theory errs by treating definite descriptions on the model of a proper name. Definite descriptions are not single words but phrases, and therefore have a richer semantic constitution.

On the singly modified naive theory, any expression other than a simple singular term is, at least in principle, capable of bifurcation of reference and information value. For example, though the information

value of a sentence is its information content, sentences might be regarded as referring to something other than their information contents. The singly modified naive theory, as defined so far, is tacit on the question of the referents of expressions other than singular terms. However, a familiar argument, due primarily to Alonzo Church and independently to Kurt Gödel, establishes that the closest theoretical analogue of singular-term reference for any expression is its extension.[1] The argument relies on three intuitive assumptions: (*a*) that a definite description 'the ϕ' refers to the only individual that satisfies the constitutive predicate (or formula) ϕ, if there is exactly one such individual, (*b*) that trivially logically equivalent referring expressions refer to the very same thing, and (*c*) that (barring such devices as quotation marks) the referent of a compound referring expression is preserved when a component referring expression is replaced by another having the very same referent. The argument is usually given for the special case of sentences, but it is easily extended to any sort of expression that has extension. Thus, for example, consider any two monadic predicates that happen to have the same extension—say, 'is a creature with a heart' and 'is a creature with kidneys' (where it is understood that a pair of monadic predicates Π and Π' are *logically equivalent* if and only if the corresponding biconditional 'Π if and only if it Π'' is a logical truth). Let us abbreviate the phrase 'the number n such that $n = 1$ if __ is a creature with a heart and $n = 0$ otherwise' by 'the degree of cordateness of __'. Thus, we define the *degree of cordateness of x* to be 1 if x is a creature with a heart and 0 otherwise. Likewise, we define the *degree of reinateness of x* to be 1 if x is a creature with kidneys and 0 otherwise. Now consider the following list of monadic predicates:

(*i*) is a creature with a heart

(*ii*) is an individual x such that the degree of cordateness of x is 1

(*iii*) is an individual x such that the degree of reinateness of x is 1

(*iv*) is a creature with kidneys.

Suppose that monadic predicates are referring expressions. By assumption *b*, predicates *i* and *ii* have the same referent. Predicate *iii* results from predicate *ii* when the (open) definite description 'the degree of reinateness of x' is substituted for the description 'the degree of cordateness of x', and by assumption *a* these are co-referential for any

value of the variable 'x'. Hence, by assumption *c*, *ii* and *iii* have the same referent. By assumption *b* again, predicates *iii* and *iv* have the same referent. Therefore, if predicates are referring expressions, then predicates *i* and *iv* have the same referent.

The same argument may be given for any pair of co-extensional monadic predicates, and indeed for any pair of co-extensional expressions. Moreover, each of the three assumptions represents fundamental principles or theorems concerning singular-term reference. Any attempt to extend the concept of reference to other sorts of expressions ought to respect these principles.

Accordingly, we modify the list of central theses of the original naive theory as follows:

Thesis I
(Declarative) sentences encode pieces of information, called *propositions*. The proposition encoded by a sentence, with respect to a given context, is its *information content* with respect to that context.

Thesis II
The information content, with respect to a given context, of a sentence is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the sentence, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices.

Thesis III′
The information value (contribution to information content), with respect to a given context *c*, of any simple singular term is its referent with respect to *c* (and the time of *c* and the world of *c*).

Thesis IV′
Any expression may be thought of as *referring*, with respect to a given context, time, and possible world, to its extension with respect to that context, time, and possible world.

Thesis V′
The information value, with respect to a given context, of a simple *n*-place first-order predicate is an *n*-place attribute (either a property or an *n*-ary relation)—ordinarily an attribute ascribed to the referents of the attached singular terms. Exceptions arise in connection with quotation marks and similar devices.

Thesis VI′
The information value, with respect to a given context, of a simple n-adic sentential connective is an attribute, ordinarily of the sorts of things that serve as referents for the operand sentences.

Thesis VII′
The information value, with respect to a given context, of a simple n-adic quantifier or second-order predicate is an n-ary attribute, ordinarily of the sorts of things that serve as referents for the operand first-order predicates.

Thesis VIII′
The information value, with respect to a given context, of a simple operator other than a predicate, a connective, or a quantifier is an appropriate attribute (for sentence-forming operators) or operation (for other types of operators)—ordinarily an attribute of or an operation on the sorts of things that serve as referents for its appropriate operands.

Thesis IX′
The information value, with respect to a given context, of a typical compound expression, if any, is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the compound expression, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices, and may arise also in connection with compound predicates. The information value, with respect to a given context, of a sentence is its information content, the encoded proposition.

2.2 *The Doubly Modified Naive Theory*

2.2.1. *Propositions and Proposition Matrices*

Although the singly modified naive theory eliminates the inconsistency built into the original naive theory, it retains a second defect of the original theory. This defect is illustrated by the following example: Suppose that at some time t^* in 1890 Frege utters the English sentence (or its German equivalent)

I am busy.

Consider the piece of information, or proposition, that Frege asserts in uttering this sentence. This is the information content of the sentence with respect to the context of Frege's uttering it. Let us call this prop-

osition 'p^*' and the context in which Frege asserts it 'c^*'. The piece of information p^* is made up of the information value of the indexical term 'I' with respect to c^* and the information value of the predicate 'am busy' with respect to c^*. According to the naive theory, these information values are Frege and the property of being busy, respectively, so p^*—the information value (= information content) of the whole sentence with respect to c^*—is a complex abstract entity made up of Frege and the property of being busy, something like the ordered couple (Frege, being busy). Let us call this complex 'Frege being busy', or '*fb*' for short. Thus, according to the naive theory, $p^* = fb$. But this cannot be correct. If *fb* is thought of as having truth value, then it is true if and when Frege is busy (if and when Frege has the property of being busy) and false if and when he is not busy. Thus, *fb* vacillates in truth value over time, becoming true whenever Frege becomes busy and false whenever he ceases being busy. (This forces a misconstrual of the intension of 'I am busy' with respect to Frege's context c^* as a two-place function which assigns to the ordered pair of both a possible world w and a time t a truth value, either truth or falsehood, according as Frege is busy in w at t or not.) But p^*, being a piece of information, has in any possible world in which Frege exists a fixed and unchanging truth value throughout Frege's entire lifetime, and never takes on the opposite truth value outside his lifetime. In this sense pieces of information are *eternal*.

Not just some; all information is eternal. The eternalness of information is central and fundamental to the very idea of a piece of information, and is part and parcel of a philosophically entrenched conception of information content. For example, Frege, identifying the cognitive information content (Erkenntniswerte) of a sentence with what he called the 'thought' (Gedanke) expressed by the sentence, wrote:

> Now is a thought changeable or is it timeless? The thought we express by the Pythagorean Theorem is surely timeless, eternal, unvarying. "But are there not thoughts which are true today but false in six months' time? The thought, for example, that the tree there is covered with green leaves, will surely be false in six months' time." No, for it is not the same thought at all. The words 'This tree is covered with green leaves' are not sufficient by themselves to constitute the expression of thought, for the time of utterance is involved as well. Without the time-specification thus given we have not a complete thought, i.e., we have no thought at all. Only a sentence with the time-specification filled out, a sentence complete in every respect, expresses a thought. But this thought, if it is true,

> is true not only today or tomorrow but timelessly. ("Thoughts," in *Logical Investigations*, pp. 27–28)

Six months from now, when the tree in question is no longer covered with green leaves, the information that the tree is *then* covered with green leaves will be misinformation; it will be false. But that information is false even now. What is true now is the information that the tree *is* covered with green leaves, i.e., the information that the tree is *now* covered with green leaves; this information is eternally true, or at least true throughout the entire lifetime of the tree in question and never false. There is no piece of information concerning the tree's foliage that is true now but will be false in six months. Similarly, if the information p^* that Frege asserts at t^* is true, it is eternally true, or at least true throughout Frege's lifetime and never false. There is no noneternal piece of information concerning Frege that vacillates in truth value as he shifts from being busy to not being busy. The complex fb is noneternal, neutral with respect to time; hence, it is not a complete piece of information, i.e., it is no piece of information at all, properly so-called.

This is not to say that the noneternal complex fb is not a semantic value of the sentence Frege utters, or that fb has nothing to do with information content. Indeed, fb is directly obtained from the sentence Frege utters in the context c^* by taking the individual associated with 'I' with respect to c^* and the property associated with 'am busy' with respect to c^*. Moreover, fb can be converted into something more like a piece of information simply by *eternalizing* it, i.e., by infusing a particular time (moment or interval) t into the complex to get a new abstract entity consisting of Frege, the property of being busy, and the particular time t. One may think of the noneternal complex fb as the matrix of the proposition p^* that Frege asserts in c^*. Each time he utters the sentence 'I am busy' Frege asserts a different proposition, expresses a different "thought," but always one having the same matrix fb. Similarly, in some cases it may be necessary to incorporate a location as well as a time in order to obtain a genuine proposition, e.g. 'It is raining' or 'It is noon'. A proposition or piece of information does not have differing truth values at different locations in the universe, any more than it has different truth values at different times. A proposition is fixed, eternal, and unvarying in truth value over both time and space.

It has been noted by William and Martha Kneale, and more recently and in more detail by Mark Richard, that this traditional conception of cognitive information content is reflected in our ordinary ascriptions of belief and other propositional attitudes.[3] As Richard points out, if what is asserted or believed were something temporally neutral or noneternal, then from the conjunction

In 1971 Mary believed (the proposition) that Nixon was president, and today she still believes that

it would be legitimate to infer

Today, Mary believes that Nixon is president.

Such an inference is an insult not only to Mary but also to the logic of English, as it is ordinarily spoken. Rather, what we may infer is

Today, Mary believes that Nixon was president in 1971.

The reason for this is that what Mary is said by the first sentence to have believed in 1971 is not the noneternal proposition matrix Nixon being president but the eternal proposition that Nixon is president (at such-and-such time) in 1971. The point is bolstered if 'know' is substituted for 'believe'.[3]

To each proposition matrix there corresponds a particular property of times, to wit, the property of being a time at which the proposition matrix is true—or, where necessary, a binary relation between times and places, to wit, the relation that obtains between a time and a place when the proposition matrix is true at that time in that place. For example, the time property corresponding to the proposition matrix *fb* is the property of being a time at which Frege is busy. It is often helpful in considering the role of proposition matrices in the semantics of sentences to think of a proposition matrix as if it were its corresponding property of times (or its corresponding relation between times and places).

2.2.2 *Information Value and Information-Value Base*

Let us call the proposition matrix that a sentence like 'I am busy' takes on with respect to a particular context *c* the *information-content base* (or, simpler, the *content-base*) of the sentence with respect to *c*. More generally, we may speak of the *information-value base* (simpler, the *value base*), with respect to a context, of a singular term, a predicate, a connective, a quantifier, etc. The value base of an expression is the entity that the expression contributes to the proposition matrix taken on by (i.e., the content base of) typical sentences containing the expression (where a "typical" sentence containing an expression does not include occurrences of such devices as quotation marks or the 'that'-operator other than those already included in the expression itself).

On the modification of the naive theory described above, the value base of a proper name, a demonstrative, or some other single-word singular term, with respect to a particular possible context *c*, would simply be the referent of the term with respect to *c*. The value base of

a simple predicate, such as 'am busy' or 'is taller than', with respect to a context c, is the attribute— property or relation—semantically associated with the predicate with respect to c, e.g. the property of being busy or the relation of being taller than. The value base of a compound expression with respect to a context c is (typically) a complex made up of the value bases of the simple parts of the compound expression with respect to c, so the value base of a sentence is just its content base. In keeping with the singly modified naive theory, the value base of a definite description, unlike that of a single-word singular term, is not simply its referent but is a complex made up partly of the property associated with the description's constitutive predicate.

Since ordinary language includes indexical expressions such as 'that tree', the value base of an expression is to be indexed to the context of utterance. This generates a new higher-level nonrelativized semantic value for an expression, on the same level as character, which is the function or rule that determines for any possible context c the value base the expression takes on with respect to c. Let us call this new semantic value the *program* of an expression. An indexical expression is precisely one that takes on different value bases with respect to different possible contexts—that is, the expression's program is not a constant function; its value base varies with the context.

The value base of an expression with respect to a context c determines a corresponding function that assigns to any time t (and location l, if necessary), an appropriate information value for the expression. (In fact, the function also determines the corresponding value base.) For example, the proposition matrix *fb*, which is the content base of 'Frege is busy' with respect to any context and also the content base of 'I am busy' with respect to any context in which Frege is the relevant agent, determines a function that assigns to any time t the information about Frege that he is busy at t. (This is the propositional function corresponding to the property of being a time at which Frege is busy.) Let us call the function from times (and locations) to information values thus determined by the value base of an expression with respect to a given context c the *schedule* of the expression with respect to c. In the special case of a single-word singular term, its schedule with respect to any context is always a constant function; however, this need not be true for other sorts of expressions, e.g. sentences. Since the information value of an expression determines its semantic intension, the value base of an expression with respect to a context c also determines a corresponding function that assigns to any time t (and location l, if necessary) the resulting intension for the expression. Let us call this function from times (and locations) to intensions the *superintension* of the expression with respect to c. Accordingly, we should speak of the

information value, and the corresponding intension, of an expression with respect to a context *c and* a time *t (and* a location *l*, if necessary).

We should also like to speak, as we already have, of the information value of an expression (e.g. the information content of a sentence) with respect to a context simpliciter, without having to speak of the information value with respect to both a context and a time (and a location). This is implicit in the notion of the character of an expression, as defined earlier. How do we get from the value base of an expression with respect to a given context to the information value with respect to the same context simpliciter without further indexing, or relativization, to a time (and location)?

In the passage quoted in subsection 2.2.1 above, Frege seems to suggest that the words making up a tensed but otherwise temporally unmodified sentence, by themselves, and even the words taken together with contextual factors that secure information values for indexical expressions such as 'that tree', at most yield only something like what we are calling a 'proposition matrix', i.e., the content base of the sentence with respect to the context of utterance, which is "not a complete thought, i.e., . . . no thought at all." He suggests further that we must rely on the very time of the context of utterance to provide a "time-specification" or "time-indication"—presumably a specification or indication of the very time itself—which supplements the words to eternalize their content base, thereby yielding a genuine piece of cognitive information, or "thought." Earlier in the same article, Frege writes:

> [It often happens that] the mere wording, which can be made permanent by writing or the gramophone, does not suffice for the expression of the thought. The present tense is [typically] used . . . in order to indicate a time. . . . If a time-indication is conveyed by the present tense one must know when the sentence was uttered in order to grasp the thought correctly. Therefore the time of utterance is part of the expression of the thought. ("Thoughts," in *Logical Investigations*, p. 10)

On Frege's view, strictly speaking, the sequence of words making up a tensed but otherwise temporally unmodified sentence like 'This tree is covered with green leaves', even together with a contextual indication of which tree is intended, does not have cognitive information content. Its information value is incomplete. Presumably, on Frege's view, the sequence of words together with a contextual indication of which tree is intended has the logico-semantic status of a predicate true of certain times—something like the predicate 'is a time at which this tree is covered with green leaves', accompanied by a pointing to the tree in question—except that the sentence may be completed by a time, serving

as a specification or indication of itself, rather than by a syntactic singular term such as 'now'. Accordingly, on Frege's theory, the information value, or "sense" (Sinn), of the sentence together with an indication of the intended tree but in isolation from any time would be a function whose values are pieces of cognitive information, or "thoughts" (Gedanke).[4] Only the sequence of words making up the sentence, together with an indication of which tree is intended and a time-indication or time-specification, as may be provided by the time of utterance itself, is "a sentence complete in every respect" and has cognitive information content.

Now, it is not necessary to view the situation by Frege's lights. Whereas Frege may prefer to speak of the cognitive thought content (Erkenntniswerte) of the words *supplemented by* both a contextual indication of which tree is intended and a "time-indication," one may speak instead (as I already have) of the information content of the sequence of words themselves *with respect to* a context of utterance and a time. The information content of 'This tree is covered with green leaves' with respect to a context c and a time t is simply the result of applying the schedule with respect to c of the sentence (sequence of words) to t. This is the singular proposition about the tree contextually indicated in c that it is covered with green leaves at t. In the general case, instead of speaking of the information value of an expression supplemented by both a contextual indication of the referents of the demonstratives or other indexicals contained therein and a "time-indication," as may be provided by the time of utterance, one may speak of the information value of the expression with respect to a context of utterance and a time (and a location, if necessary). Still, Frege's conception strongly suggests a way of constructing a singly indexed notion of the information value of an expression with respect to (or supplemented by) a context of utterance c simpliciter—without further relativization to (or supplementation by) a time (and a location)—in terms of the doubly (or triply) indexed locution: The information value of an expression with respect to a context c (simpliciter) is definable as the information value of the expression with respect to both c and the very time of c (and the very location of c, if necessary).

In particular, then, the information content of a sentence with respect to a given context of utterance c is its information content with respect to c and the time of c (and the location of c, if necessary). Consequently, any present-tensed but otherwise temporally unmodified sentence encodes different information with respect to different contexts of utterance (simpliciter). For example, Frege's sentence 'This tree is covered with green leaves' encodes different information with respect to different times of utterance, even when pointing to the same tree. Uttered now,

it encodes the information about the tree in question that it is covered with green leaves, i.e., that it is now covered with green leaves. Uttered six months from now, it encodes the information about the tree that it is then covered with green leaves. This is precisely the phenomenon Frege noted and attempted to capture with his remark that the time of utterance completes the sentence as part of the expression of its thought content.

Let us call this latest version of the naive theory the *doubly modified naive theory* (abbreviated as simply the *modified naive theory*). The doubly modified naive theory is the singly modified naive theory modified further to accommodate the eternalness of information value.

It follows from our definition of the singly indexed notion of the information value of an expression with respect to a context simpliciter that the program of an expression fully determines the expression's character, since, given any context *c*, the program fully determines the resulting schedule, which together with the time (and location, if necessary) of *c* fully determines the resulting information value. From this it follows that the program of an expression also determines the expression's contour, as defined earlier. Within the framework of the modified naive theory, the meaning of an expression is better identified with its program, rather than with its character. This allows one to distinguish pairs of expressions like 'the U.S. president' and 'the present U.S. president' as having different meanings, despite their sharing the same (or nearly the same) character. More accurate, the program of an expression is the primary component of what is ordinarily called the 'meaning' of the expression, though an expression's meaning may have additional components that supplement the program.[8]

The original and the singly modified naive theory recognize three distinct levels of semantic value. The three primary semantic values are *extension*, *information value* (misconstrued as possibly noneternal), and *character*. In addition, these theories admit two subordinate semantic values. On the same level as, and fully determined by, information value is *intension* (misconstrued as a two-place function from possible worlds and times); on the same level as, and fully determined by, character is *contour*. The various semantic values on the original or the singly modified naive theory, and their levels and interrelations, are diagrammed in figure 1. (Of course, these are not the only semantic values available on the naive theory, but they are the important ones.)

The modified naive theory's notion of the value base of an expression with respect to a given context, and the resulting notion of the program of an expression, impose a fourth level of semantic value, intermediate between the level occupied by Kaplan's notion of the character of an expression and the level of information value. In fact, the introduction

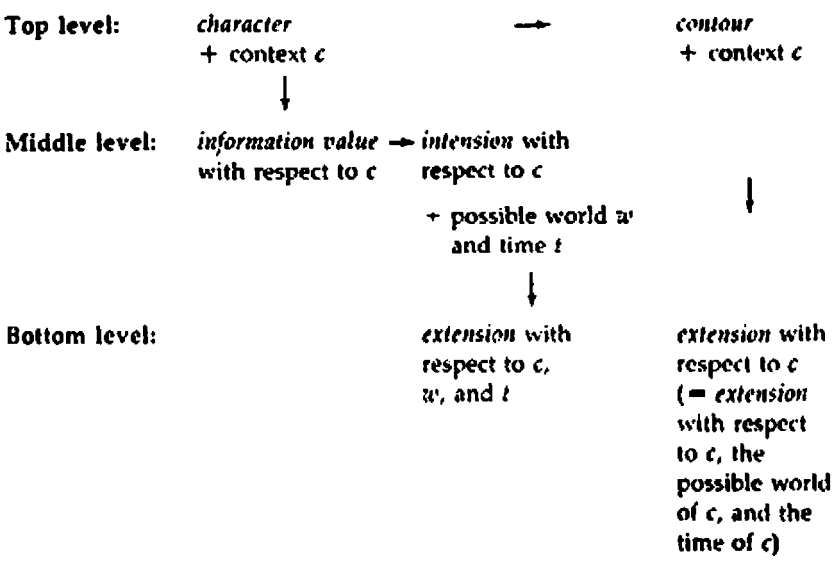


Figure 1
Semantic values on the naive theory.

of the notion of value base reduces character to the status of a subordinate semantic value. The four primary semantic values, from the bottom up, are *extension*, *information value* (construed now as necessarily eternal), *information-value base*, and *program*. In addition, there are a number of subordinate semantic values. Besides intension (construed now as a one-place function from possible worlds), character, and contour, there are *schedule* and *superintension*, both of which are on the same level as, and fully determined by, value base. The various semantic values on this modification of the original naive theory, and their levels and interrelations, are diagrammed in figure 2. (Notice that figure 1 is virtually embedded within figure 2, as its right half.)

On the modified naive theory, the extension of an expression with respect to a given context of utterance (simpliciter, without further relativization to a time, a place, or a possible world) is the result of applying the intension of the expression with respect to that context—which, in turn, is the result of applying the superintension of the expression with respect to that context to the very time of the context—to the very possible world of the context. Thus, for example, the referent of a singular term—say, 'the U.S. president's actual wife'—with respect to a particular context of utterance c is semantically determined in a sequence of steps. First, the program of the expression is extracted from

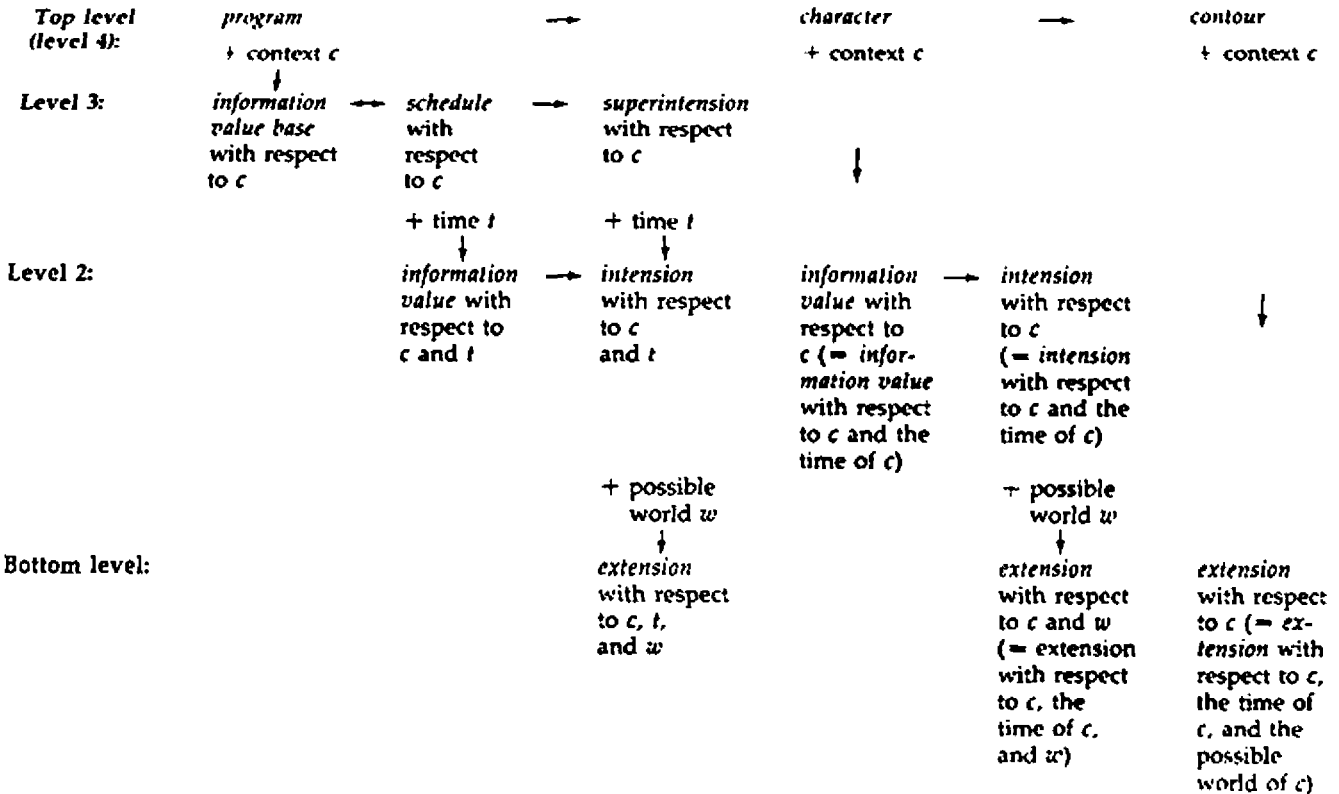


Figure 2
Semantic values on the modified naive theory.

its meaning. This program is then applied to the context c to yield the time-neutral value base of the expression with respect to c. This value base yields the schedule of the expression with respect to c, which assigns to any time t the information value of the expression with respect to both c and t. This schedule is applied to the very time of c itself to give the eternal information value of the expression with respect to c (simpliciter). This information value, in turn, yields the expression's intension with respect to c, which assigns to any possible world w the extension of the expression with respect to c and w. The expression's extension with respect to any context c' and possible world w' is the individual who is the wife in the possible world of c' at the time of c' of the individual who is the U.S. president in w' at the time of c'. Finally, this intension is applied to the very possible world of c itself to yield the individual who is the wife in the possible world of c at the time of c of the individual who is the U.S. president in the possible world of c at the time of c. Eureka!

2.2.3 *Tense versus Indexicality*

It may appear that I have been spinning out semantic values in excess of what is needed. We needed a singly indexed notion of the information value of an expression with respect to a context, and as a special case, a notion of the information content of a sentence with respect to a context. This led to the original and singly modified naive theories' identification of meaning with character. In the special case of a single-word singular term, what I am calling its value base with respect to a context c is the very same thing as its information value with respect to c, so the program of a single-word singular term is just its character. The only thing that prevents this from holding also for a sentence like 'I am busy' is that its content base with respect to a context is neutral with respect to time, whereas its information content with respect to the same context is eternal, somehow incorporating the time (and location, if necessary) of the context. It may seem, then, that, in the case of a sentence or phrase, what I am calling its 'value base' with respect to a context c is just its information value with respect to c but for the deletion of the time of c (and the location of c), so that the information content of a sentence with respect to a context c is made up of the information values (= value bases) with respect to c of its simple information-valued parts *plus* the time (and location) of c. However, if the rule of information-content composition is that information contents are constructed from the information values of the simple information-valued components *together with* the time (and location, if necessary) of utterance, then why bother mentioning those partially constructed pieces of information I have been calling 'proposition ma-

trices'? Singling out content bases as separate semantic values generates the doubly indexed notion of the information content of a sentence with respect to both a context c and a time t, and thereby the non-relativized higher-level notion of program. But what is the point of this doubly indexed notion, and of the resulting notion of program? Are we not interested only in the case where the time t is the time of the context of utterance c (and where the location l is the location of the context c)? Why separate out the time as an independent semantic parameter that may vary independent of the context of utterance? The character of a sentence seems to be meaning enough for the sentence. Semantic theorists heretofore have gotten along fine by indexing the notion of information content once, and only once, to the context of utterance, without relativizing further and independently to times. For example, in discussing the phenomenon of tense, Frege considers also various indexicals—'today', 'yesterday', 'here', 'there', and 'I'—and suggests a uniform treatment for sentences involving either tense or indexicals: "In all such cases the mere wording, as it can be preserved in writing, is not the complete expression of the thought; the knowledge of certain conditions accompanying the utterance, which are used as a means of expressing the thought, is needed for us to grasp the thought correctly. Pointing the finger, hand gestures, glances may belong here too." ("Thoughts," in *Logical Investigations*, pp. 10–11)

Following Frege, it would seem that we can handle the phenomena of tense and indexicality together in one fell swoop, with tense as a special case of indexicality, by simply relativizing the notion of information value once and for all to the complete context of utterance—including the speaker and his or her accompanying pointings, hand gestures, and glances as well as the time and location of the utterance. Any aspect of the complete context of utterance may conceivably form "part of the expression of the thought" or contribute to the information content. Once information content is relativized to the complete context, including the time of utterance, gestures, and so on, there seems to be no need to relativize further and independently to times.

It has become well known since the middle of the 1970s that the phenomenon of tense cannot be fully assimilated to temporal indexicality, and that the presence of indexical temporal operators necessitates "double indexing," i.e., relativization of the extensions of expressions—the reference of a singular term, the truth value of a sentence, the class of application of a predicate (or better, the semantic characteristic function of a predicate), etc.—to utterance times independent of the relativization to times already required by the presence of tense or other temporal operators.[6] (Something similar is true in the presence of an indexical modal operator such as 'it is actually the case that' and in

the presence of indexical locational operators such as 'it is the case here that'.) Though it is less often noted, it is equally important that double indexing to contexts and times (or triple indexing to contexts, times, and locations, if necessary) is required at the level of information value (e.g. information content) as well as at the level of extension (e.g. truth value).[7] For illustration, consider first the sentence

At t^*, I believed that Frege was busy.

By the ordinary laws of temporal semantics, this sentence is true with respect to a context of utterance c if and only if the sentence 'I believe that Frege is busy' is true with respect to both c and the time t^*. This, in turn, is so if and only if the binary predicate 'believe' applies with respect to c and t^* to the ordered pair of the referent of 'I' with respect to c and t^* and the referent of the 'that'-clause 'that Frege is busy' with respect to c and t^*. Hence, the displayed sentence is true with respect to c if and only if the agent of c believes at t^* the piece of information that is the referent of the 'that'-clause 'that Frege is busy' with respect to c and t^*. What piece of information does the 'that'-clause refer to with respect to c and t^*? The information content of its operand sentence 'Frege is busy', of course. (See the discussion of the 'that'-operator in the introduction.) But which proposition is that? If information content is to be singly indexed to context alone, it would seem to be the information content of 'Frege is busy' with respect to c. This is the proposition that Frege is busy at t, where t is the time of c. However, this yields the wrong truth condition for the displayed sentence. This would be the correct truth condition for the sentence 'At t^*, I believed that Frege would be busy now'. The displayed sentence ascribes, with respect to c, a belief at t^* that Frege is busy at t^*. Assuming that information content is singly indexed to context alone, we are apparently forced to construe the 'that'-operator in such a way that a 'that'-clause 'that S' refers with respect to a context c and a time t' not to the information content of S with respect to c but to the information content of S with respect to a (typically different) context c' which is exactly like c in every aspect (agent, location, etc.) except that its time is t'. (The contexts c and c' would be the same if and only if t' were the time of c.) This yields the desired result that the displayed sentence is true if and only if the agent of c believes at t^* the information that Frege is busy at t^*.

This account appears to yield exactly the right results until we consider a sentence that embeds an indexical temporal operator within the 'that'-operator and embeds the result within another temporal operator. Consider the following:

> When the Republicans next regain the U.S. presidency, Jones will believe that the present U.S. president is the best of all the former U.S. presidents.

This sentence is true with respect to a context c if and only if the time (assuming there is such) at which the Republicans next regain the presidency after the time t of c—let us call this time 't''—is such that Jones believes at t' the piece of information referred to by the 'that'-clause 'that the present U.S. president is the best of all the former U.S. presidents' with respect to c and t'. On the singly indexed account of information content, the displayed sentence comes out true if and only if Jones believes at t' that the U.S. president at t' is the best of all the U.S. presidents before t'. But this is the wrong truth condition for the displayed sentence. In fact, it is the correct truth condition for the nonindexical sentence obtained by deleting the word 'present'. The displayed sentence ascribes, with respect to c, a belief that the U.S. president at t (the time of c) is the best of all the U.S. presidents prior to t'. In order to obtain this result, the 'that'-clause 'that the present U.S. president is the best of all the former U.S. presidents' must be taken as referring with respect to c and t' to the proposition that the U.S. president at t is the best of all the U.S. presidents prior to t' (or to some proposition trivially equivalent to this). This cannot be accommodated by a singly indexed account, and it requires seeing information content as doubly indexed: to the original context c (so that the ascribed belief concerns the U.S. president at t rather than t') and the time t' (so that the ascribed belief concerns the class of U.S. presidents before t' rather than those before t).

2.2.4 *Temporal Operators*

The example just considered illustrates the need for the double indexing of information content that is generated by the modified naive theory's notion of the content base of a sentence. In addition to this, there is an important semantic function for the content base of a sentence that cannot be fulfilled by its information content. To see this, it is important to look more closely at the semantics of temporal operators.

Consider the temporal connective 'sometimes', which attaches to a sentence S to form the sentence 'Sometimes S'. An appropriate extension for this operator would be a function from some aspect of the operand sentence S to a truth value. What aspect of the operand sentence S?

Two sorts of operators are very familiar to philosophers of language. An *extensional* operator is one that operates on the extensions of its operands, in the sense that an appropriate extension for the operator itself would be a function from *extensions* appropriate to the operands

(as opposed to some other aspect of the operands) to extensions appropriate to the compounds formed by attaching the operator to an appropriate operand. An extensional sentential connective (such as 'not' or 'if . . . , then – – –') is one that is truth functional; an appropriate extension would be a function from (*n*-tuples of) truth values to truth values, and hence an appropriate information value would be an attribute (property or relation) of truth values—for example, the property of being falsehood, or the following relation: Either u is falsehood or v is truth. An *intensional* or *modal* operator is one that operates on the intensions of its operands. An appropriate extension for a modal connective like 'it is necessarily the case that' or 'if it were the case that . . . , then it would be the case that – – –' would be a function from (*n*-tuples of) sentence intensions (functions from possible worlds to truth values) or propositions to truth values, and an appropriate information value would be an attribute of intensions or propositions—for example, the property of being a necessary truth.

Now, is the 'sometimes' operator extensional? Certainly not. With respect to my actual present context, the sentences 'It is cloudy' and '2 + 2 = 5' are equally false, though 'Sometimes it is cloudy' is true whereas 'Sometimes 2 + 2 = 5' is false. Thus, the 'sometimes' operator is not truth functional, and hence not extensional. Nor is the 'sometimes' operator intensional, in the above sense. With respect to my actual present context, the two sentences 'The senior senator from California is a Republican' and 'The present senior senator from California is a Republican' have precisely the same intension—indeed, they have (very nearly) the same information content—though 'Sometimes the senior senator from California is a Republican', on the relevant reading (the Russellian secondary occurrence or small scope reading), is true whereas 'Sometimes the present senior senator from California is a Republican', on either of its two readings (Russellian small scope vs. large scope), is false. Thus, 'sometimes' is neither an extensional operator nor an intensional or information content operator. What, then, is it?

In order to obtain the correct results, one must regard a sentential temporal operator such as 'sometimes' as operating on some aspect of its operand sentence that is fixed relative to a context of utterance (so as to give a correct treatment of temporally modified indexical sentences like 'Sometimes the present senior senator from California is a Republican') but whose truth value typically varies with respect to time (so that it makes sense to say that it is *sometimes* true, or true at such-and-such time). On the original and singly modified naive theories' three-tiered array of semantic values (as diagrammed in the right half of figure 2 in subsection 2.2.2 above), once it is acknowledged that information content is eternal, there simply is no such semantic value

of a sentence. Nothing that is fixed relative to a context is also time-sensitive in the required way. In order to find an appropriate semantic value for temporal operators such as 'sometimes' to operate on, one must posit a level of semantic value intermediate between character and information content. This strongly suggests that the objects of sentential temporal operators—the things operated on by sentential temporal operators such as 'sometimes'—are something like proposition matrices, or perhaps sentence superintensions. The 'sometimes' operator is neither an extensional operator nor an intensional (i.e., modal) operator, nor is it even a contour operator (since sentence contours are context-sensitive entities, and would thus yield incorrect results for temporally modified indexical sentences like 'Sometimes the present senior senator from California is a Republican'). Instead, 'sometimes' is a superintensional operator. That is, an appropriate extension for 'sometimes' with respect to a context c, a time t, and a possible world w would be a function from the superintension (or, equivalently, from the schedule or content base) of its operand sentence (with respect to c) to a truth value—namely, the function that assigns truth to a proposition matrix (or to its corresponding schedule or superintension) if its value for at least one time (the resulting proposition or sentence intension) itself yields truth for the world w, and which otherwise assigns falsehood to the proposition matrix (or the corresponding schedule of superintension). In general, temporal operators—such as 'sometimes', tense operators (including complex ones such as present perfect and future perfect), indexical temporal operators (e.g. 'present'), and even nonindexical specific time indicators (e.g. 'on December 24, 1996' + future tense or 'when Frege wrote "Thoughts" ' + past tense)—may all be seen as superintensional operators.

A sentence of the form 'Sometimes S' may be regarded as encoding, with respect to a given context c, information concerning the content base of the operand sentence S with respect to c. For example, the sentence 'Sometimes I am busy' contains, with respect to Frege's context c^* (or any other context in which Frege is the agent), the information about the proposition matrix *fb* that it is sometimes true. Accordingly, an appropriate information value for a temporal operator such as 'sometimes' would be a property of proposition matrices—in this case, the property of being true at some time(s). It is in this way that temporal operators like 'sometimes' provide a place for proposition matrices in temporal semantics and thereby generate a doubly indexed notion of the information value of an expression (e.g. the information content of 'I am busy' or of 'This tree is covered with green leaves') with respect to both a context c and a time t that may be other than the time of c. Just as it is the information content of its operand that a modal operator

says something about (e.g. that it is a necessary truth), so it is the information content base of its operand that a temporal operator says something about.

2.2.5 Predicates and Quantifiers

An important point about predicates, quantifiers, and certain other operators emerges from the four-tiered modified naive theory, and from the distinction between information value and value base in particular: The value base of a predicate such as 'is busy' or 'is taller than', with respect to a given context of utterance c, is an attribute, i.e., a property or relation. This, together with a time t, determines the information value of the predicate with respect to c and t. In turn, the information value of a predicate with respect to c and t, together with a possible world w, determines the extension of the predicate with respect to c, t, and w. It follows that the information value of a predicate such as 'is busy' with respect to a context c and a time t is not just the property of being busy (or anything similar, such as the function that assigns to any individual x the proposition matrix, x being busy). The property of being busy together with a possible world w cannot determine the extension of 'is busy' with respect to both the world w and the time t. The property of being busy together with a possible world w determines only the class of (possible) individuals who are busy *at some time* in w, or, at most, the function that assigns to any time t the class of (possible) individuals who are busy at t in w. The information value of 'is busy' with respect to a given time t must be such as to determine for any possible world w the class of (possible) individuals who are busy *at the given time* t in w. Only some sort of complex consisting of the property of being busy *together with the given time* t is such as to determine for any possible world w the extension of 'is busy' with respect to both w and t. Thus, the information value of 'is busy' with respect to a given time t is not merely the property of being busy but a complex consisting of this property and the time t. This, it may be assumed, is a *temporally indexed attribute*—in this case, the temporally indexed property of *being busy at* t. Similarly, the information value of 'is taller than' with respect to a time t is the temporally indexed binary relation of *being taller than at* t, which is made up of the nonindexed binary relation of being taller than and the time t. In general, the information value of a predicate with respect to a time t (and a location l, if necessary) is not the same attribute as the value base of the predicate but is the temporally (and, if necessary, spatially) indexed attribute that results from taking the value base of the predicate together with the time t (and the location l, if necessary). This heretofore unrecognized

fact about the information values of predicates allows us to retain, at least as a sort of general guide or rule of thumb, the principle that the information value of a compound expression, such as a sentence or phrase, is a complex made up solely and entirely of the information values of the information-valued components that make up the compound. In particular, the information content of 'I am busy' with respect to a context of utterance c may be thought of as made up of the agent of the context c and the property of being busy at t, where t is the time of c. There is no need to introduce the time t as a third and separate component; it is already built into the information value of the predicate.

Exactly analogous remarks apply to quantifiers, other second-order predicates, the definite description operator 'the', and certain other operators. Accordingly, the list of central theses of the (doubly) modified naive theory is the following:

Thesis I′
(Declarative) sentences encode pieces of information, called *propositions*. The proposition encoded by a sentence, with respect to a given context and time, is its *information content* with respect to that context and that time.

Thesis II′
The information content, with respect to a given context and time, of a sentence is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the sentence, typically the simple (non-compound) component expressions. Exceptions arise in connection with quotation marks and similar devices.

Thesis III″
The information value (contribution to information content), with respect to a given context c and time t, of any simple singular term is its referent with respect to c and t (and the world of c).

Thesis IV′
Any expression may be thought of as *referring*, with respect to a given context, time, and possible world, to its extension with respect to that context, time, and possible world.

Thesis V″
The information value, with respect to a given context c and time t, of a simple n-place first-order predicate is an n-place attribute,

ordinarily an attribute temporally indexed to t (either a temporally indexed property or a temporally indexed n-ary relation)—ascribed to the referents of the attached singular terms. Exceptions arise in connection with quotation marks and similar devices.

Thesis VI″
The information value, with respect to a given context and time, of a simple n-adic sentential connective is an attribute, ordinarily of the sorts of things that serve as referents for the operand sentences.

Thesis VII″
The information value, with respect to a given context c and time t, of a simple n-adic quantifier or second-order predicate is an n-ary attribute, ordinarily an attribute temporally indexed to t of the sorts of things that serve as referents for the operand first-order predicates.

Thesis VIII″
The information value, with respect to a given context and time, of a simple operator other than a predicate, a connective, or a quantifier is an appropriate attribute (for sentence-forming operators) or operation (for other types of operators), ordinarily an attribute of or an operation on the sorts of things that serve as referents for its appropriate operands.

Thesis IX″
The information value, with respect to a given context and time, of a typical compound expression, if any, is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the compound expression, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices, and may arise also in connection with compound predicates. The information value, with respect to a given context and time, of a sentence is its information content, the encoded proposition.

Since the information value of an expression with respect to a context c simpliciter is the information value with respect to both c and the time of c (and the location of c, if necessary), it follows that the information value of a typical predicate with respect to a context c simpliciter varies with the context c—whether or not the predicate

indexical (such as one possible reading of 'is current', as in 'is a current journal issue'), and hence even if it is not, like 'is busy' or 'is taller than'. It is this previously unnoticed feature of predicates that accounts for the fact that a nonindexical temporally unmodified sentence, e.g. 'Frege is busy', takes on not only different truth values but also different information contents when uttered at different times, even though the sentence is not indexical. It is also this feature of predicates that accounts for the fact that certain noneternal (i.e., temporally nonrigid) definite descriptions, such as 'the U.S. president', take on not only different referents but also different information values when uttered at different times, though the description is not indexical. Recall that the distinctive feature of an indexical like 'I' or 'the present U.S. president' is that it takes on different information-value bases in different contexts. The predicate 'is busy', the definite description 'the U.S. president', and the sentence 'The U.S. president is busy' all retain the same value base in all contexts. Their information value varies with the context, but not their value base.

The account presented here of the information values of temporal operators as properties of proposition matrices (or other value bases) makes for an important but usually unrecognized class of exceptions to the general principle that the information value of a compound expression is made up of the information values of its information-valued components. Where T is a monadic temporal sentential operator, e.g. 'sometimes' + present tense or 'on July 4, 1968' + past tense, the information content of the result of applying T to a sentence S, with respect to a context c, is made up of the information value of T with respect to c together with the information-content *base* of S with respect to c, rather than the information content of S itself. In general, if T is a temporal operator, the information value with respect to a context c of the result of applying T to an expression is a complex made up of the information value of T with respect to c and the value base, rather than the information value, of the operand expression with respect to c. Ordinarily, the information value of an expression containing as a part the result of applying a temporal operator T to an operand expression is made up, in part, of the value base of the operand expression rather than its information value. (For complete accuracy, the notion of information value with respect to a context, a time, and a location, for a language L, should be defined recursively over the complexity of expressions of L.)[9]

The modified naive theory retains all the cogency and intuitive appeal of the original naive theory while correcting for the inconsistency engendered by the latter's extreme naiveté, and while accommodating the eternalness of information. The modified naive theory combines

some of the sophistication and power of Fregean semantic theory, which allows for bifurcation of reference and information value, with the natural simplicity of the original view, and is thus even more compelling than the original naive theory. Unfortunately, Frege's question arises even on the modified naive theory—in the special case where the a and b in '$a = b$' are proper names or other single-word singular terms.

What makes Frege's challenging question a puzzle is that it poses a seemingly insurmountable problem for a seemingly indisputable account of the nature and structure of information content—or, more accurate, for what would be a seemingly indisputable account were it not so vigorously disputed. And as I hope to show, as with many other philosophical puzzles, no solution proposed so far is entirely satisfactory, and most are clearly unsatisfactory.

Chapter 3

The Theories of Russell and Frege

3.1 *Russell*

Contemporary philosophers who have shown considerable sympathy with something like the naive theory, or some modification thereof, include Keith Donnellan, David Kaplan, Saul Kripke, Ruth Barcan Marcus, and John Perry. I suspect that the list of closet naive-theory sympathizers is a long and distinguished one.

Historically, the staunchest and foremost champion of something very much like the singly modified naive theory is Bertrand Russell. In fact, firm adherence to some variant or other of the naive theory may be the only consistent and unwavering theme throughout Russell's philosophical career, the one thesis immune from revision in Russell's thought.[1] Russell handled Frege's Puzzle (and the other puzzles associated with the naive theory) by introducing an ingenious wrinkle in the theory. His idea was both brilliant and far-reaching: Though the naive theory is fundamentally correct in identifying the information value of a name with its referent and in countenancing singular propositions as the basic units of information, one must not be misled by the surface structure of an atomic subject-predicate sentence as to its propositional content. In the standard case, the proposition assigned to such sentences must be more complex than it would first appear. What appears to be a genuine name or deictic term is usually, according to Russell, not a singular term at all but a "denoting phrase"—a certain sort of expression that signals the presence of a quantificational construction.[2] The sentence 'Socrates is wise' does not assert anything directly about Socrates at all; Socrates does not "occur as a constituent" of the relevant proposition. Instead, the sentence asserts something about a certain pair of properties, or, more accurate, about a certain pair of *propositional functions*. The word 'Socrates' is not a genuine name of Socrates. Instead, it is semantically correlated with a certain propositional function that is in some sense *definitional* of 'Socrates' (say, the propositional function *being a snub-nosed Athenian philosopher sentenced to death for his views and social activities*). Let us call this

propositional function '*Socrateity*'. Then the sentence expresses something equivalent to the ascription of a certain complex relation between *Socrateity* and the propositional function *wisdom*: *Socrateity* is uniquely instantiated, and, in addition, *Socrateity* and *wisdom* are co-instantiated. (An individual x is said to *instantiate* a propositional function F if the proposition obtained by applying F to x is true, and a pair of propositional functions are said to be *co-instantiated* if there is some individual that instantiates both.) For example, on one interpretation of Russell's theory, the sentence encodes the information about the complex propositional function *being a unique Socratizer who is wise* that it is instantiated (where a *Socratizer* is anything that instantiates *Socrateity*).

The apparatus of Russell's sophisticated revamping of the naive theory begins with the singular propositions (more accurate, the proposition matrices) of the original naive theory, but it is an additional tenet of Russell's theory that singular propositions concerning individuals are never entertained, except perhaps in the special case of individuals of intimate epistemic aquaintance, and even then they are entertained only very briefly and never communicated. Instead, we deal primarily with propositions directly concerning certain intensional entities: propositional functions. In particular, an apparent identity sentence '$a = b$', in the ordinary case, will express not a singular proposition about the referents of the names a and b but a considerably more complex proposition (equivalent to one) about a certain pair of propositional functions F_a and F_b that they are co-uniquely-instantiated, whereas the sentence '$a = a$' will express (a proposition equivalent to) the much weaker proposition about the propositional function F_a that it is uniquely instantiated.

3.2 *Frege*

If Russell is cast in the role of Ptolemy, adding ingenious epicycles to the conventional wisdom of the ages in order to make it fit the recalcitrant facts, then Frege must be cast in the coveted role of Copernicus. The analogy is especially apt in the way it depicts the nature of the chasm between Russell and Frege.

There are a great many important similarities and points of contact between the theories of Russell and Frege.[3] Both theories see ordinary proper names (such as 'Socrates') and ordinary uses of indexicals and demonstratives as providing not the object but a sort of conceptual representation employing certain intensional entities ordinarily correlated with predicates, as a component of the piece of information encoded by containing sentences. But any account that is content

assimilate the theories of Russell and Frege and to ignore the differences misses an important and dramatic element of the whole picture.

Faced with a recalcitrant phenomenon (the informativeness of identity sentences using two names), Frege makes a revolutionary proposal: One must not merely modify the naive picture. Instead, the naive theory must be scrapped altogether, to be replaced by a new philosophy of semantics in which the second tier of semantic value (information value) is compounded not of elements of the first tier (reference) but of a special realm of entities. On Frege's theory, any meaningful expression, whether a sentence component or a complete sentence, semantically refers to (designates, stands for) something, if anything, as its referent (Bedeutung), but it only does so by semantically *expressing* something else: its *sense* (Sinn). The sense of an expression is a purely conceptual representation, and the referent of the expression is whatever uniquely fits the representation.[4] In the terminology of Alonzo Church, the referent of an expression is whatever its sense uniquely *determines*. The sense of the name 'Socrates' is the name's purely conceptual "mode of presentation" of the individual Socrates, and the name refers to this individual, rather than someone else, by virtue of the fact that Socrates is the only individual who fits, or is determined by, the semantically associated conceptual representation. The sense of an expression secures the expression's referent. Moreover, an expression's sense is its information value. The sense of an expression is, thus, a semantically associated purely conceptual representation that forms part of the cognitive information content of sentences containing the expression, and the referent of the expression is whatever happens to fit this representation uniquely. Nothing counts as the sense of an expression, properly so-called, unless it is, all at once, the expression's semantically associated purely conceptual "mode of presentation," the mechanism that secures the expression's referent, and the expression's information value. In claiming that a name such as 'Socrates' has sense for a particular user of the name, Frege is identifying the information value of the name for a particular user with the purely conceptual content the user associates with the name. Frege held that the sense of a compound expression, such as a sentence, is a product of the senses of its parts, and, similarly, that the referent of a compound is a function of the referents of its parts. (See note 4 to chapter 1.) He often spoke of the senses of the parts of a compound as themselves parts of the sense of the compound. The referent of a sentence like 'Socrates is wise' is simply its truth value, either truth or falsehood, whereas its sense is its cognitive information content (Erkenntniswerte) and is purely conceptual in nature. Frege called these special senses 'thoughts' (Gedanke). It is not Socrates himself but a conceptual representation—the sense

of 'Socrates'—that goes into the information or thought that Socrates is wise. Information cannot involve concrete individuals (or even sense data) as constituents, but must consist solely of conceptual entities. Fregean "thoughts" correspond roughly to a special subclass of Russellian propositions, namely the purely general propositions.[5] From the point of view of the Fregean theory, in placing Socrates as a component of the information that Socrates is wise, the naive theory (and any modification thereof) rests on a confusion of extreme dimensions, and no amount of tinkering or adjusting can adequately correct for it. (Compare the Copernican view of the Ptolemaic model of the solar system.)

Frege's theory of sense solves Frege's Puzzle by pointing out that, though $\ulcorner a = b \urcorner$ and $\ulcorner a = a \urcorner$ ascribe the same relation to the same pair of objects, this does not mean that they encode the same information, that they have the same "cognitive value." The first sentence can be informative even though the second is not because the information contained in the first sentence is made up, in part, of the senses of both the names a and b, whereas the information contained in the second sentence is not. The sentences therefore convey different pieces of information.

I emphasize the differences between the theories of Russell and Frege because unless one fully appreciates the radical nature of this divergence one misses the central point of Frege's theory altogether. One example of the sort of misunderstanding (or at least misdescription) I am talking about may be found in a passage from Leonard Linsky, a contemporary champion of Frege's theory of sense and an otherwise lucid expositor of Frege's views. Expounding Frege's solution to the puzzle, he writes:

> . . . once [someone] does discover (either by himself or from others) the identity [of Hesperus and Phosphorus], his discovery will be that the [sense] associated by him with 'Hesperus' and that associated by him with 'Phosphorus' pick out the same object, Venus. . . . His discovery is that two (individual) concepts are concepts of the same object. Thus what [he] discovers is not merely a fact about words or names. 'Hesperus = Phosphorus' does not mean that 'Hesperus' denotes the same object as 'Phosphorus'. (Linsky, *Names and Descriptions*, p. 72)

It is true that, on Frege's theory, the discovery or information ("thought") that Hesperus is Phosphorus is not information about the names 'Hesperus' and 'Phosphorus'. It is also true that for Frege the sense of a name is relevant to—indeed, partly constitutive of—cognitive information content. But the information that Hesperus is Phosphorus is not about senses any more than it is about names. It is about a massive physical object, the planet Venus. On Frege's theory, the in-

formation that Hesperus is Phosphorus is one thought and the information that the individual concept *Hesperus* determines the same object as the individual concept *Phosphorus* is another. (The information that the concepts *Hesperus* and *Phosphorus* both determine Venus is yet a third thought.) The first piece of information is made up of the individual concepts *Hesperus* and *Phosphorus* (the senses of the names 'Hesperus' and 'Phosphorus'), which are concepts of (i.e., determine) a certain planet, whereas the second is made up of concepts of these concepts. Contrary to the impression left by Linsky, for Frege the sentence 'Hesperus = Phosphorus' does not mean (i.e., encode the information) that the individual concepts *Hesperus* and *Phosphorus* determine the same object any more than it means that the names 'Hesperus' and 'Phosphorus' refer to the same object. It means simply that the *objects* Hesperus and Phosphorus *are* the same object. Linsky's characterization of Frege's account is at best only half correct. Linsky is correct in pointing out that the concepts expressed by 'Hesperus' and 'Phosphorus' are what, according to Frege, go into the information that Hesperus is Phosphorus, but he grossly misrepresents Frege's theory when he concludes that the information is therefore information about these concepts (to wit, that they determine the same object). In fact, in this Linsky is lapsing back into the naive theory's account of the nature of information. In effect, he mistakes the Fregean thought that Hesperus is Phosphorus, which (assuming it exists) is *made up* partly of the individual concepts *Hesperus* and *Phosphorus*, for a singular proposition about these concepts—a proposition that is, coincidentally, composed of some of the same things as the Fregean thought. But the singular proposition is something of a sort that Frege would vigorously reject as having nothing to do with the cognitive information content of the identity sentence—unless, say, one regards the singular proposition about Venus that it is it as a concoction that mathematically represents a special equivalence class of which the thought content of the identity sentence is an element. Even then, none of the other members of this infinite class are, for Frege, the cognitive information content of the identity sentence. A singular proposition, qua information content, is a gadget resulting from a theory that Frege believes to be sheer confusion. The sort of theory that Linsky appears to attribute to Frege—one according to which the information content of 'Hesperus = Phosphorus' is a singular proposition about a pair of intensional entities correlated with the names 'Hesperus' and 'Phosphorus', respectively, to the effect that these intensional entities single out the same object—is perfectly coherent, and is indeed a plausible proposal for dealing with Frege's Puzzle. But the theory is not Frege's; it is Russell's. It is the sort of theory that Frege is aiming to debunk. The central point of Frege's

theory of cognitive value is that, contrary to the naive account, a piece of cognitive information is not made up of the thing or things it is "about", in the ordinary and relevant sense of 'about'. For Frege, the thought that Socrates is wise, though about Socrates, does not have Socrates as a component part, and likewise the thought that the individual concept *Socrates* determines a wise man does not have the concept *Socrates* as a component part. There are for Frege countless ways of conceiving any object, countless purely conceptual modes of presenting the object to the mind's grasp, and each of these can go into the makeup of a different thought. The thought is about an object only by virtue of containing a sense that determines the object, not by containing the object itself.[7]

Conventional philosophical wisdom since Russell has tended to favor Frege's solution to Frege's Puzzle over Russell's. Like Ptolemy's mature model of the solar system, Russell's sophisticated revamping of the naive theory is fundamentally a fine-tuned adjustment of an older picture, an ad hoc variation on an older theme. Ingenious though it is, it retains few if any of the merits enumerated above in connection with the original or the modified naive theory, on which the information value of a proper name or a demonstrative, as ordinarily used, is identified with the referent. Whatever else may be said in favor of Russell's account, the advantages of the original view do not apply. Still less do they apply to Frege's theory, but Frege's theory may boast important advantages of its own. Russell's theory, with its commitment to and emphasis on the epistemological primacy of sense data, has fallen upon hard times. The idea that my belief that this table is wooden is only indirectly about the table, and directly about a private experience, is implausible. The idea that we seldom if ever communicate the precise content of our thought to others is perhaps even less plausible. Recent philosophy, forgoing Russell's variant of the naive theory but finding Frege's alternative unsatisfactory on several counts, has revived some of the central elements of the modified naive theory. Such is the way of the theory of direct reference.[8] (Compare contemporary philosophy's finding a way to interpret the naive theory in such a way as to make it acceptable with contemporary physics' finding a way to interpret Ptolemy's account in terms of relativity theory so that it is, in a sense, acceptable.)

The naive theory once held Frege in its grip. In section 8 of "Begriffsschrift" he wrote that, in the usual case, singular terms "are merely representatives of their content [referent] so that every combination into which they enter expresses only a relation between their respective contents [referents]." Fully aware of the challenging question to which this account gives rise, he made an exception for the special case of

identity statements. In these contexts, the singular terms flanking the identity predicate '= ' were held to refer to themselves. In Frege's "Begriffsschrift," the sentence '$\overline{a} = \overline{b}$' is supposed to express the metatheoretic singular proposition about the singular terms *a* and *b* that they are co-referential.

By the time he came to write "Über Sinn und Bedeutung," Frege found reason to reject this analysis of identity statements. This is all for the good, since the "Begriffsschrift" account, taken as an analysis of natural-language identity statements, is surely mistaken. A number of objections have been raised in the literature,[9] some better than others, but one serious objection that is not usually noted (at least not in quite this form) follows directly from the first point made in section 1.1 above concerning what Frege's Puzzle is a puzzle about. I argued that the central elements of Frege's Puzzle have nothing special to do with the identity relation or the identity predicate, since different versions of the very same puzzle can arise in connection with any predicate whatsoever. Consequently, no purported solution that proposes a reinterpretation or a new analysis for just the special case of identity statements can remove the general problem. We are still left with no answer to the challenging questions that remain in connection with 'Shakespeare wrote *Timon of Athens*' and with 'Hesperus is a planet if Phosphorus is'. If there is a general solution that provides answers to these questions, it should extend straightforwardly to the special case of 'Hesperus is Phosphorus'. The general problem concerns the analysis of the sort of information that is semantically contained in declarative sentences, the feature of sentences that accounts for their informativeness or their uninformativeness. An adequate solution to the puzzle must address this issue directly. The "Begriffsschrift" account is thus not only mistaken but irrelevant. Unless it is only part of a sweeping proposal for reinterpreting declarative sentences generally, as with the theories of Hobbes and Mill, and not just a proposal concerning identity sentences, it does not even speak to the issue raised by the puzzle.[10]

In contrast, Frege's later solution invoking the sense-referent distinction speaks directly to the issue. His own objection to his earlier "Begriffsschrift" proposal, however, was not that it fails to solve the general problem. It was that the "Begriffsschrift" account radically misrepresents the nature of the information or the fact conveyed in a genuine identity sentence '$\overline{a} = \overline{b}$'. Frege objected that the fact that the two singular terms *a* and *b* are co-referential is "arbitrary." "One cannot," he wrote, "forbid anyone to take any arbitrarily produced event or object as a sign for anything. Accordingly, a sentence $a = \overline{b}$' would no longer concern the thing itself, but only our mode of designation; we would express no genuine knowledge therein." From the viewpoint

of the theory of "Über Sinn und Bedeutung," of all of the objections that have been raised against the "Begriffsschrift" analysis of identity statements, this one of Frege's must surely be seen as one of the weakest. The objection constitutes a great irony in Frege's philosophy. Frege claims that the fact that two names—say, 'Hesperus' and 'Phosphorus'—happen to name the same thing is an uninteresting accident of the use of language, a result of arbitrary linguistic convention, and is irrelevant to the subject matter—in this case, astronomy—determined by the object so named, whereas the fact that the objects Hesperus and Phosphorus are the same thing is an interesting fact of astronomy and is independent of human decision or convention. What Frege failed to notice is that this claim is categorically denied by the very theory of sense that he uses this argument to motivate! On that theory, the two names function semantically in a manner very similar to that of a definite description, even if the name is not strictly synonymous with any description available from natural language.[11] For on Frege's theory of sense, the referent of a name like 'Hesperus' is secured by means of a semantically associated conceptual mode of presentation—perhaps the representational content of a description of the form 'the first heavenly body visible in such-and-such location at dusk'—and likewise the referent of 'Phosphorus' is secured by means of a semantically associated representational content—perhaps that of 'the last heavenly body visible in so-and-so location at dawn'. Now, the fact that these descriptions refer to the same thing is indeed due in part to certain "arbitrary" accidents of the English language. It is due in part to the semantical facts that 'the' + NP refers to the unique thing satisfying NP, that 'heavenly body' means what it does, and so on. No population can be forbidden from using these constructions in some way other than this, if they are going to be that way about it. But just as certainly, the fact that the two descriptions are co-referential is not due solely to these "arbitrary" accidents of linguistic usage. It is due also to the nonarbitrary fact, independent of human activity, that some one heavenly body is visible in such-and-such location at dusk later than any other heavenly body and also visible in so-and-so location at dawn earlier than any other heavenly body. Were it not for this fact, which is not about linguistic usage but about the existence of a certain kind of heavenly body, the two descriptions mentioned above simply would not be co-referential. The fact that the descriptions are co-referential is not just a result of human decision, convention, or of human activity in general. It is also the result of a certain celestial state of affairs. By the same token, if proper names have a Fregean sense, the fact that two names like 'Hesperus' and 'Phosphorus' refer to the same thing is not just a result of arbitrary linguistic convention but is due also to a fact con-

cerning the subject matter determined by the object referred to. On Frege's theory of sense, it is solely a result of semantic stipulation or linguistic convention or decision which sense is attached to a particular proper name, at least in certain cases. One cannot be forbidden from using a name with whatever sense one chooses, but that is the extent of stipulatory input into the referent of the name. Once a particular sense is attached to a name, the matter of which object this sense determines is decided by the extralinguistic facts, independent of further semantic stipulation. Human decision may fix the sense for the name, but once this is done we simply sit back and this sense independently determines its object, namely whatever object uniquely fits the conceptual mode of presentation decided upon. Of course, if it is discovered that this is not the intended object, the sense of the name can be changed by a further decision; this, however, does not alter the fact that, whatever sense is ultimately attached to the name, the referent will depend on which object happens to fit the sense uniquely (not to mention the fact that the mistake may never be discovered). Thus, on the orthodox theory of proper names, the matter of which object a name names is not solely the result of human decision but is due also to independent facts concerning which object uniquely satisfies the particular mode of presentation attached to the name.

Why did Frege suppose and insist that the fact that two singular terms are co-referential is merely the result of arbitrary linguistic convention or usage, if his theory precludes this? My conjecture is that, even on the very brink of the announcement of his revolutionary proposal for solving the puzzle, Frege was subconsciously still under the powerful, seductive spell of something like the modified naive theory.[12] On the modified naive theory, there is no reason to suppose that the fact that a name names a particular thing is anything but the result of human linguistic activity, and there is every reason to suppose that the fact is just that. For the modified naive theory does not accord proper names the sort of semantic autonomy that is characteristic of definite descriptions and other phrases, whereby the expression relies on a semantically contained mode of presentation to secure an extension, whatever happens to fit the conceptual mode of presentation. If proper names do not have this sort of mechanism for securing a referent, how then do they secure a referent unless, ultimately, by something like semantic stipulation or usage? The fact that a proper name names a particular thing must be due to speakers' using the name to refer to that particular thing, or intending that the name should be so used, or something of the sort. Indeed, the fact that the name names what it does is in some significant sense *constituted* by this sort of linguistic activity. This linguistic activity may be causally related to certain ex-

traterrestrial states of affairs, but the linguistic activity is itself wholly mundane.

The sense theory's claim that the matter of which object a name refers to is due in part to the extralinguistic fact of which object uniquely fits the particular associated concept yields an implausible account of what makes it true that a name names what it does, at least in the standard sort of case where someone in authority has conferred the name on someone or something. Rather, Frege's contratheoretical observation that one wholly and simply *decides* to let some expression be a name for someone or something in particular seems essentially correct. Given the appropriate conventions and institutions, and within certain constraints thereby laid down, the parents of a newborn child simply stipulate what the child's name will be. When my wife befriends yet another stray or abandoned cat, it is usually left up to me to name it. When I do, I do not first assign some sort of conceptual description (say, *the calico cat that Eileen just adopted, whichever cat that turns out to be*) and then allow this concept to probe the universe seeking whatever fits it uniquely (crossing my fingers in hopes that I got things right). I choose a name, and I begin referring to the cat by that name. I look the cat straight in the eye and I say 'You will be Sonya'. I have thereby stipulated that 'Sonya' will be the name for *this very cat*, irrespective of her color or breed or how she became a member of the household. If instead I had decided to bind a certain concept to the name 'Sonya' and to rely on this concept to provide a referent, as Frege's theory would have it, Frege would simply be mistaken in insisting that the circumstance of 'Sonya' naming the relevant cat obtains by stipulation. What the name 'Sonya' refers to, if anything, would also depend on which cat, if any, happens to satisfy certain conditions uniquely. But here Frege is correct; it is his theory that is mistaken.[13] Ironically, while arguing in favor of his theory of sense, Frege explicitly acknowledged that the matter of which object a proper name names is due entirely to linguistic usage or human decision or activity—a fact his theory is forced to deny. Frege's theory of sense thus undermines one of his primary arguments in its favor. More important, Frege's modified-naive-theory-motivated observation concerning the trivial, stipulatory character of the reference of proper names constitutes an important argument against the very theory he is about to announce in the paragraph following the observation. Surely there can be no greater testimony to the cogency of the modified naive theory.

4

The Structure of Frege's Puzzle

4.1 Compositionality

I have claimed that Frege's Puzzle concerns the nature and structure of pieces of information (the sort of information semantically contained in a declarative sentence), and that an adequate solution must address this issue directly. It is important for this purpose to focus on the principles and assumptions involved in the derivation of Frege's Puzzle.

Preliminary investigation into the nature and structure of pieces of information uncovered that a piece of information is a complex abstract entity whose components are the information values of the components of a sentence that contains the information (modulo the qualifications mentioned in note 4 to chapter 1). There are two components of the information that Socrates is wise: what is had in common between the information that Socrates is wise and the information that Socrates is snub-nosed, and what is had in common between the information that Socrates is wise and the information that Plato is wise. It is natural to suppose that the first component is precisely the individual whom that information is about, i.e., the man Socrates. Frege's Puzzle challenges this natural idea by proposing two purportedly distinct pieces of information that have the very same predicative component and are about the very same individual. The implicit assumption is a principle of compositionality for pieces of information: If pieces of information are complex abstract entities, and two pieces of information p and q having the same structure and mode of composition are numerically distinct, then there must be some component of one that is not a component of the other; otherwise p and q would be one and the very same piece of information. (Compare the principle of extensionality for classes or sets.)

This compositionality principle for pieces of information might be challenged. Complex entities having the very same components and mode of composition cannot always be identified with one another. The clipboard on which I am now writing has the very same component molecules as the matter that now constitutes it, but, for familiar philo-

sophical reasons, the clipboard is not identical with its present matter. The clipboard came into existence long after its present matter did, and it will cease to exist long before its present matter does (if the matter ever ceases to exist). Moreover, strictly speaking, the clipboard is constituted by different (albeit largely overlapping) matter at different times, and is only briefly constituted by its present molecules, though the present matter is forever constituted by these very molecules. Similarly, to use an example due to Richard Sharvy, the Supreme Court of the United States has the very same membership as the set of its present justices, but the Court and the set of its present justices are distinct complex entities, since the Court changes its membership over time whereas no set can change its membership.[1] Even complex entities of the very same kind having the same constituents and mode of composition cannot always be identified. Different ad hoc committees within a university department can coincide exactly in membership though they remain different committees with different functions and responsibilities.

In contrast with these examples, it would seem that pieces of information do obey the principle of compositionality implicit in Frege's Puzzle. For each of the complex entities mentioned above as violators of a corresponding compositionality principle, there is some significant aspect of the entity, some crucial feature of it, that differentiates it from any distinct entity composed of the very same constituents in the very same way. The Supreme Court and the set of its present justices differ in their flexibility with respect to change in membership. Any two distinct ad hoc committees differ in at least some of their functions or purposes. But pieces of information having the very same structure and components, combined in the very same way, cannot change in constitution, and they fulfill the same purposes and perform the same functions. In any event, if two pieces of information, p and q, are composed of the very same components in the very same way but are distinct, it would seem that there must also be some important aspect in which they differ, some significant property had by p and not by q or vice versa. This, however, raises the same challenging question posed by Frege, or at least a philosophically important question similar to Frege's original question: What in the world is this mysterious feature or aspect of pieces of information in which two pieces of information composed of the same components in the same way can yet differ? Even if the principle of compositionality for pieces of information fails, some variant of Frege's Puzzle remains a pressing philosophical problem for semantic theory.

4.2 *Frege's Law*

In order to produce two distinct pieces of information that are about the same individuals and that have the same predicative component, Frege offers a pair of declarative sentences involving the same predicate but different singular terms for the same object and argues that these sentences must be seen as containing different pieces of information. To this end, Frege's Puzzle, in its original form, tacitly invokes the following principle concerning information content:

> If a declarative sentence *S* has the very same cognitive information content (Erkenntniswerte) as a declarative sentence *S'*, then *S* is informative ("contains an extension of our knowledge") if and only if *S'* is (does).

I shall call this principle *Frege's Law*. It is an exceedingly plausible principle connecting the concepts of *information content* and *informativeness*. Still, it might be thought that it is precisely the unquestioning acceptance of this principle that is the source of the puzzle. It might even be argued that the puzzle should be recast as a reductio ad absurdum of the principle. "What independent reason can there be," one might ask, "for holding this principle to be true? In fact, isn't it clear that the informativeness or uninformativeness of a sentence depends on more factors than just the information content of the sentence, so that two sentences having the same content may yet differ in their informativeness?"[2]

This line of attack against Frege's Puzzle is sorely mistaken. Given the sense of 'informative' that is relevant to the puzzle, Frege's Law is unassailable. Properly understood, Frege's Law should be seen as a special instance of Leibniz's Law, the Indiscernibility of Identicals. This is because, on a proper understanding of 'informative', the informativeness or uninformativeness (a posteriority or a priority, etc.) of a sentence is a *derivative* semantic property of the sentence, one that the sentence has only by virtue of encoding the information that it does encode. That is, to say that a sentence, on a particular occasion of use, is (as the term is used in the context of Frege's Puzzle) *informative* (or that it is *a posteriori*) is to say something about the *information content* of the sentence: It is to say that the information content is not somehow already given, or that the content is nontrivial, or that it is knowable only by recourse to experience and not merely by reflection on the concepts involved, or that it is an "extension of our knowledge," or something along these lines. There is some such property *P* of pieces of information such that a sentence is informative, in the sense relevant to Frege's Puzzle, if and only if its information content has the property *P*.

Of course, there are other senses of 'informative' on which even a trivial identity statement may be described as "informative". For example, if you do not speak a word of French but you have it on good authority that Jean-Paul's next inscription will be of a true French sentence, and you observe Jean-Paul then write the words 'Cicéron est identique à Cicéron', the sentence in question, on this occasion of use, may be said to be "informative" on several counts. By way of its inscription, you are given a great deal of nontrivial information; you are thereby given that a certain sequence of marks is a meaningful and grammatical expression of French, that it is in fact a French sentence, and that it is a true sentence. If you also know even a minimum about the grammar of Romance or Indo-European languages, and you know that 'Cicéron' is a name, you are also thereby given the information that the words 'est identique à' probably signify some relation in French, a relation that the relevant person called 'Cicéron' in French bears to himself. However, all this is quite irrelevant to Frege's Puzzle.

It is extremely important in dealing with Frege's Puzzle and related philosophical problems to distinguish the notion of the *information content* of a sentence on a particular occasion of its use from the notion of the *information imparted* by the particular utterance of the sentence. The first is a semantic notion, the second a pragmatic notion. Failure to make this distinction has led many a well-meaning philosopher astray. I have already discussed the notion of semantically encoded information at some length in the previous chapter. In claiming that it is a basic function of sentences to encode information, I invoke the notion of semantically encoded information. To illustrate the quite different notion of pragmatically imparted information, it is best to begin with a nonlinguistic and uncontroversial example. Consider some of the ways in which one might receive or learn the information that Smith has a cold. One way, of course, is for someone (perhaps Smith) to produce with assertive intent a conventional symbol that semantically encodes that information; for example, Smith may utter the sentence 'I have a cold' in conversation. Under certain circumstances, another way to learn that Smith has a cold—one not involving language—is simply to observe Smith sneeze and then blow his nose. In this sense, Smith's blowing his nose imparts, or can impart, the information that he has a cold. Though the blowing of a nose may thus impart certain information, it would be utterly ridiculous to suppose that nose blowing has any semantic content. One can imagine a society in which blowing one's nose is a linguistic gesture—a move in the language game—much like shaking one's head 'no' is in our society; fortunately, however, we do not live in such a society. In our society, nose blowing has no semantic significance whatsoever. It is an entirely nonlinguistic act.

Now, just as Smith's nose blowing may impart the information that Smith has a cold, without itself having any semantic attributes and hence without semantically encoding any information, so any observable event typically imparts some information to the astute observer—hence the saying "Actions speak louder than words." Utterances are no exception. In uttering a sentence, one produces a symbol that semantically encodes a piece of information, and in so doing one performs an action (indeed, several actions) that, like any other action, may impart information in the nonsemantic way that even nose blowing may impart information. Of course, typically the information semantically encoded by a sentence will be pragmatically imparted by utterances of the sentence. But the two notions may diverge and often do. In addition to (sometimes instead of) the information semantically encoded by a sentence, an utterance of the sentence may impart further information concerning the speaker's beliefs, intentions, and attitudes, information concerning the very form of words chosen, or other extraneous information. The further information thus imparted can often be of greater significance than the information actually encoded by the sentence itself. Such is the case with Jean-Paul's inscription of 'Cicéron est identique à Cicéron'. In this sense, even utterances can "speak louder than words." In particular, one piece of information typically imparted by the utterance of a sentence S is the information that S is true with respect to the context of the utterance. It is rarely the case, however, that a sentence semantically encodes the information about itself that it is true (or, for that matter, that it is not true—such is the stuff of which paradoxes are made).

Frege himself was aware of the distinction between semantically encoded and pragmatically imparted information. Using his word 'thought' (Gedanke) for what I am calling 'information', Frege explicitly drew the distinction, or something very similar to it, in a section entitled "Separating a Thought from its Trappings" of an essay entitled "Logic," estimated to have been composed in 1897:

> . . . we have to make a distinction between the thoughts that are expressed and those which the speaker leads others to take as true although he does not express them. If a commander conceals his weakness from the enemy by making his troops keep changing their uniforms, he is not telling a lie; for he is not expressing any thoughts, although his actions are calculated to induce thoughts in others. And we find the same thing in the case of speech itself, as when one gives a special tone to the voice or chooses special words. (in *Posthumous Writings*, ed. Hermes et al., at p. 140)

Frege's Puzzle concerns only the information content of Jean-Paul's sentence—the nature and structure of the information semantically

contained in or encoded by the sentence with respect to the particular context of use—and not the information pragmatically imparted by the particular utterance. When Frege claims that sentences of the form $\ulcorner a = a \urcorner$ are a priori and do not "contain very valuable extensions of our knowledge," and are in this respect different from sentences of the form $\ulcorner a = b \urcorner$, there is no question but that he is concerned only with the "thought expressed" by this form of sentence, i.e. its information content, and not with the unexpressed "thoughts" that the utterance "leads us to take as true." The information content of Jean-Paul's sentence is utterly trivial. It is in this essentially semantic sense of 'informative', having to do with the character of the information encoded by a sentence, that this French sentence is quite definitely uninformative. Its information content is a *given*, and does not "extend our knowledge."

To take another example due to Carnap, consider the numerical equation '5 = V', using both the Arabic and the Roman numeral for five.[3] To someone familiar with one but not both of these numeral systems, an inscription of this equation pragmatically imparts nontrivial information, e.g. the information concerning one of the numerals that it is a numeral for the number five. But the information semantically encoded by the equation is precisely the same as that encoded by '5 = 5'. This is an instance of the trivial law of reflexivity of equality. The encoded information is not a "valuable extension of knowledge," or anything of the sort. In the relevant sense, the equation is utterly uninformative. A similar situation obtains with respect to sentences like 'Opthalmologists are oculists' and 'Alienists are psychiatrists'. To someone unfamiliar with the grammatical subject term but familiar with the grammatical predicate term, an utterance or inscription of one of these sentences pragmatically imparts nontrivial linguistic information concerning the meaning of the grammatical subject term, though the semantically encoded information is utterly trivial. Indeed, it is just this feature of these sentences—the fact that their semantic information content is trivial—that suits them to the task of conveying the meanings of 'ophthalmologist' and 'alienist'. This is unlike the examples that give rise to Frege's Puzzle (e.g. 'Hesperus is Phosphorus'), in which we are to suppose that the audience has complete mastery of both terms and finds the utterance or inscription informative nevertheless.

Properly understood, then, Frege's Law is not merely a plausible principle connecting the concepts of information content and informativeness, or even a fundamental law of semantics. It is a truth of logic. Hence, it is no solution to the puzzle to challenge Frege's Law.

4.3 *Challenging Questions*

I have argued that it is no help to deny the information-compositionality principle implicit in Frege's Puzzle, and that, properly understood, there is no denying Frege's Law. Given the further premise that there are pairs of sentences of the forms ϕ_a and ϕ_b that differ in informativeness even though a and b are co-referential proper names, demonstratives, single-word indexical singular terms, or any combination thereof, we have all of the makings of a refutation of the modified naive theory and the consequent challenging philosophical questions for semantic theory. Suppose that there are such pairs of sentences, say the pair 'Hesperus is Hesperus' and 'Hesperus is Phosphorus'. By Frege's Law and the compositionality principle, it follows that atomic pieces of information, such as the information that Hesperus is Phosphorus, are not always singular propositions. In fact, at least one of the sentences 'Hesperus is Hesperus' and 'Hesperus is Phosphorus' must encode a piece of information that is not the same thing as the singular proposition about the planet Venus that it is it. In particular, either the name 'Hesperus' or the name 'Phosphorus', or both, contributes as its information value something other than its referent (the planet Venus) to the relevant piece of information.[4] And, for reasons of symmetry, evidently both names must have something other than the referent as information value. Here the challenging questions arise. What is this mysterious thing that is the information value of the name? Given Frege's Law and the compositionality principle, the information value of each name cannot be simply the referent of the name, as the modified naive theory would have it. So what is it? That it is just the information value of the name is, of course, no answer. We seek illumination, not labels. What sort of thing is the information value of a name or of a demonstrative in use? Is it identifiable with something that is specifiable independent of the concept of information? Is it perhaps a *concept*? (Frege's answer is something like this.) Is it a complex consisting of an object and a concept? Is it a linguistic entity—perhaps the name itself—or a complex consisting of an object and the name? Each of these proposals has been entertained at one time or another by at least one philosopher (Nathan Salmon). Unfortunately, none of them works any better than the modified naive theory's claim that the information value of a name, a demonstrative, or some other single-word singular term is simply its referent. I shall attempt to show this in the following chapter.

5

A Budget of Nonsolutions to Frege's Puzzle

Is there anything to be learned from Frege's Puzzle? What, if anything, does the puzzle show about proper names, demonstratives, or other single-word singular terms?

5.1 Conceptual Theories

5.1.1. The Orthodox Theory versus the Theory of Direct Reference

Frege, Russell, and their followers use Frege's Puzzle as an important argument—often the main argument—establishing the inadequacy of both the original and the modified naive theories and demonstrating the need to see ordinary uses of proper names, demonstratives, and certain other indexicals as contributing to the information content of containing sentences not their referent (their semantically correlated individual) but instead a conceptual representation of the referent (semantically correlated individual) built from the same sort of intensional entities that serve as the information values of predicates.[1] This conceptual representation is called upon to do double duty: In addition to being the information value of the term, it also serves as the semantic mechanism by which the referent of the term (the semantically correlated individual) is secured and semantically determined, in the straightforward sense that the referent (semantically correlated individual) of the term, with respect to a possible world and a time, is supposed to be whoever or whatever uniquely fits the representation in that world at that time.[2] On the theories of both Frege and Russell, the conceptual content of a name or indexical a in ordinary use semantically determines the truth conditions (or the semantic superintension, i.e., the semantically correlated function from times to functions from possible worlds to truth values) of its containing sentence ϕ_a, in the sense that the truth value of the sentence (with respect to a possible world w and a time t) is determined by semantics alone to be truth if the unique object fitting the conceptual content of a (with respect to w and t) satisfies ϕ (with respect to w and t).

A more complete statement of the theory held in common by Frege and Russell can be given by drawing some distinctions. Let us say that an expression *a*, as used in a particular possible context, is *descriptional* if there is a set of properties or concepts semantically associated with *a* in such a way as to generate a semantic relation, which may be called 'denotation' or 'reference' and which correlates with *a* (with respect to semantic parameters such as a possible world *w* and a time *t*) whoever or whatever uniquely has or fits all (or at least sufficiently many) of these properties or concepts (in *w* at *t*), if there is a unique such individual, and nothing otherwise. A descriptional term is one that refers or "denotes" by way of properties or concepts. It is a term that expresses a way of conceiving something, and its referent or "denotation" (with respect to a possible world and a time) is secured indirectly by means of this conceptual content. Definite descriptions, such as 'the author of *The Republic*', are the paradigm descriptional expressions. A nondescriptional singular term is one whose reference is not semantically mediated by associated conceptual content. The paradigm of a nondescriptional singular term is the individual variable. An individual variable is a singular term that does not refer or denote simpliciter, but refers under an assignment of values to individual variables. As pointed out in the introduction, the referent of a variable (with respect to a possible world and a time) under such an assignment is semantically determined directly by the assignment, and not by extracting a conceptual "mode of presentation" from the variable.

Frege and Russell held that proper names, demonstratives, and other indexical singular terms, as used in ordinary contexts, are descriptional. But they held more than this. On their view, if the name 'Saint Anne' is analyzable as 'the mother of Mary', it must be in some sense analyzable even further, since the name 'Mary' is also supposed to be descriptional. But even 'the mother of the mother of Jesus' must be in this sense further analyzable, in view of the occurrence of the name 'Jesus', and so on.

Let *a* be a nondescriptional singular term referring to Socrates. Then the definite description ⌜the wife of $\overline{a}$⌝, though descriptional, is not thoroughly so. The concept expressed is not one like that of being married to the philosopher who held that such-and-such. Rather, it is an intrinsically relational concept involving Socrates directly as a constituent: the concept of being *his* wife. We may say that the description is only *relationally descriptional*, and that it is *descriptional relative to* Socrates. A *thoroughly descriptional* term, then, is one that is descriptional but not relationally descriptional.[3]

The orthodox theory, as advocated by Frege and Russell, is that proper names, demonstratives, and other indexical singular terms ('you', 'here', etc.), as used in a particular possible context, are either thoroughly

descriptional or descriptional relative only to items of "direct acquaintance," such as sensations and visual images.

Frege held the strong version of this theory that proper names, demonstratives, and other indexical singular terms, as used in a particular context, are all thoroughly descriptional: Only if a term is thoroughly descriptional can there be something that counts as a genuine Fregean sense for the term. The reason for this is that, as I noted in section 3.2 above, the Fregean conception of sense is a compilation or conflation of three distinct linguistic attributes. First, the sense of an expression is a purely conceptual mode of presentation. Individuals that are not themselves senses, e.g. persons and their sensations, cannot form part of a genuine Fregean sense. Second, the sense of a singular term is the mechanism by which its referent is secured and semantically determined. Third, the sense of an expression is its information value. Nothing counts as the sense of a term, as Frege intended the notion, unless it is all three at once. It is supposed that the purely conceptual content of any singular term is also its information value, which also secures its referent. This three-way identification constitutes a strong theoretical claim. A descriptional singular term is precisely one whose mode of securing a referent is its descriptive content, which also serves as its information value. Only if the term is thoroughly descriptional, however, can this be identified with a *purely* conceptual (or a *purely qualitatively* descriptive) content. Even a Russellian term that is descriptional relative only to items of direct acquaintance (if there are any such terms) does not, strictly speaking, have a genuine Fregean sense. (See chapter 3, note 4.)

Since the mid 1960s, the orthodox theory, as advocated by Frege, Russell, and their followers, has been forcefully challenged by a number of philosophers, most notably Keith Donnellan, David Kaplan, Saul Kripke, and Hilary Putnam. It has been effectively demonstrated by example and argument that the conceptual or descriptive content of a proper name, demonstrative, or single-word indexical singular term will often befit nothing whatsoever, will consequently fail to determine the correct referent and to yield the correct truth conditions, and in some cases may even determine the wrong person or object altogether. These considerations form the starting point of the *theory of direct reference*, according to which proper names, demonstratives, and single-word indexical singular terms, in ordinary use, are nondescriptional.[4] These considerations, however, do not show, and are not put forward to show, that such singular terms, in use, are devoid of conceptual or descriptive content, since they clearly and obviously do evoke concepts in the minds of speakers. No one can seriously maintain that the mind typically draws a complete blank whenever 'Socrates,' 'Shakespeare',

or the name of a familiar acquaintance is used. What kind of a language user would a person whose mind did draw a blank be? The disagreement between the direct-reference theorists and the orthodox theorists is not over the *existence* of conceptual or descriptive content, but rather over the alleged semantic role of conceptual or descriptive content in securing a referent and contributing to truth conditions (with respect to semantic parameters). Direct-reference theorists emphasize the role of nonconceptual, contextual factors in the securing of a referent. However, it is perfectly compatible with the theory of direct reference, as propounded by Kripke in *Naming and Necessity* and by others elsewhere, that the information value of a name is, at least partly, its conceptual or descriptive content. Perhaps the Fregeans are correct, then, in locating the information value of a name, demonstrative, or other single-word indexical singular term, on a given occasion of its use, in its purely conceptual content on that occasion. Granted that the full-blown descriptional theory is untenable as a theory of reference, should this minimal component of the orthodox theory concerning information value be accepted nonetheless?

No. There are arguments that show that the information value of a name or an indexical is not its conceptual or descriptive content. In fact, some of the arguments offered by direct-reference theorists to show that the conceptual content of a name or an indexical is not what secures its referent can, with slight modification, also be made to show that the conceptual content is not the information value either.

5.1.2 The Twin-Earth Argument

One compelling argument against identifying the information value of a name with its purely conceptual content can be extracted from Hilary Putnam's twin-earth thought experiment.[5] The argument depends on two plausible assumptions. First, it assumes that one's (purely psychological) state of consciousness determines which (purely conceptual or purely qualitative) concepts one is grasping, in the sense that if person A is in the very same (purely psychological) state of consciousness as person B, then, for any (purely conceptual or purely qualitative) concept c, A grasps c if and only if B also grasps c.[6] Second, the argument assumes that the information component that corresponds to the individual that a given piece of information is about determines that individual, in the sense that, if a piece of information p is information about an individual x and the component of p corresponding to x is also (appropriately) a component of a piece of information q, then q is also information concerning x. For example, on this assumption, if the information value of the name 'Socrates' is appropriately part of a piece of information p, then p is information concerning Socrates. Now, sup-

pose that in a far corner of the universe there is a planet on which there is a perfect duplicate of a particular earthly woman. Each lives a life on her own planet qualitatively identical to the other's. Even their mental streams of consciousness are qualitatively identical. Moreover, each has a husband named 'Hubert', and the two Huberts are dead ringers for one another except that the earthly Hubert weighs exactly 165 pounds whereas his alien counterpart weighs exactly 165.000000001 pounds. Now, suppose that both wives simultaneously utter, assertively and sincerely, the string of symbols 'Hubert weighs exactly 165 pounds' in conversation, each talking about her own husband. The speakers are in exactly the same (purely psychological) state of consciousness. In fact, their very brain matter is in exactly the same configuration, molecule for molecule. Hence, by the first assumption, the purely conceptual content that each associates with her use of the name 'Hubert' is exactly the same. But the information encoded by the sentence uttered, as used on these two occasions, is different. This is evident because the information asserted by the earthly woman concerns her husband and is true whereas the information asserted by the alien woman concerns her husband and is strictly false. Hence, by the second assumption, the information value of the name 'Hubert' as used by the two women is different. The purely conceptual content is the same, but the information value is different. It follows that the information value of a name cannot be simply its purely conceptual content.

5.1.3 Further Arguments

The twin-earth argument shows that, contrary to the Fregean theory, purely conceptual content cannot be the whole of information value for proper names. Perhaps, then, the associated descriptive or conceptual content of a name is only a part of the information value. It might be proposed that the information value of a name is constituted partly by associated descriptive or conceptual content and partly by the referent.[7]

Could the information value of a name be something like the ordered couple of the referent together with the conceptual content? Unfortunately, this proposal fares no better than the original identification of information value with purely conceptual content. One immediate difficulty with locating the information value of a name even only partly in its conceptual content, whether purely conceptual or not, arises from the subjectivity of conceptual content. The conceptual content of a name, as used with reference to a certain person or thing, varies widely among those who have learned to use the name with that reference. The concept attached to a name by users well acquainted with the bearer of the name may be quite different from that attached by users who know the bearer only in passing, and both of these are

quite different from the concept attached by a user who only knows of the bearer but does not know the bearer personally. Moreover, conceptual content varies considerably among users of each sort. As Frege noted, some people may think of Aristotle only as the pupil of Plato who taught Alexander the Great, while others may think of Aristotle only as the teacher of Alexander the Great born in Stagira. If conceptual content is information value, or even just a part of information value, then the information encoded by a sentence containing a name will vary from person to person exactly as much as the conceptual content each attaches to the name. This idea clashes sharply with the original, natural idea of a sentence—e.g. 'Socrates is wise'—encoding a single piece of information (the information that Socrates is wise). The sentence 'Socrates is wise', as used with reference to the famous snub-nosed philosopher, encodes the same information for you as for me. More accurate, the sentence, so understood, encodes a single piece of information, period. It does not encode a piece of information *for* someone. (See chapter 1, note 2.) The encoding relation between sentences and pieces of information is a nonsubjective semantic attribute that is every bit as objective as the semantic attributes of truth and falsehood. It is not to be relativized subjectively to persons or their idiosyncratic associations.

This observation requires a couple of caveats. First, the encoding relation must be relativized to a particular *type of use* of the sentence. The sentence 'Aristotle wrote *The Metaphysics*', as used with reference to the Stagirite philosopher, encodes an uncontroversial piece of information. As used with reference to the late shipping magnate Aristotle Onassis, it encodes a piece of misinformation. This type of relativization occurs also with the attributes of truth and falsehood. Relativization of the encoding relation to types of use is necessary to ensure a definite, unambiguous reference for any names contained in the sentence (among other things). It is relativization to a particular assignment of a referent to the name. One might even prefer to say that it is relativization to a particular *name* (as opposed to other names with the same spelling and pronunciation but a different referent). In either case, it is not the same thing as relativization to a particular conceptual content. Relativization to a type of use is a necessary precondition for semantic attribution; conceptual associations are irrelevant to semantic attribution. Relativization to a particular type of use, or to a particular name-with-referent, does not result in a plurality of information contents or truth values, one "for" this reader of the sentence and another "for" that. Once a particular type of reference use is fixed upon (e.g. use with reference to the Stagirite philosopher), the sentence with that use unambiguously encodes a single piece of information with a single truth

value—one for all, or, better, one not "for" any. Second, because the conceptual content of a name varies from user to user, the sentence "Aristotle wrote *The Metaphysics*', as used with reference to the Stagirite philosopher, may well convey different pieces of information to different users. But recall the distinction between semantically encoded and pragmatically imparted information. An utterance of any sentence typically imparts more information to the audience than merely the information semantically encoded. What information is imparted depends, in part, on the idiosyncratic conceptual associations made by the listener or reader. The sentence "Aristotle wrote *The Metaphysics*' may impart some information to one reader and different (perhaps overlapping) information to another, but it encodes a single piece of information, and, if all goes well, that encoded information is part (though only part) of the information imparted. As we have seen, Frege's Puzzle and the attendant notion of information value concern only the notion of semantically encoded information. Idiosyncratic associations are beside the point.

A further argument against the proposal to locate the information value of a proper name even partly in its descriptive or conceptual content comes directly from the modal and epistemological arguments advanced by Kripke for the theory of direct reference. Suppose that the descriptive content one associates with the name 'Shakespeare'—one's concept of Shakespeare, one might say—includes some particular property as a central or critical element, say the authorship of *Romeo and Juliet*. If the proposed theory of information value is correct, the information encoded by the sentence 'If Shakespeare exists, then he wrote *Romeo and Juliet*' must be both necessarily true and knowable a priori. However, it is neither necessary nor a priori. It might have come to pass that Shakespeare elected to become a lawyer instead of a writer and dramatist. Furthermore, it is easy to imagine circumstances in which it is discovered that, contrary to popular belief, Shakespeare did not write *Romeo and Juliet*. Since this possibility is not automatically precluded by reflection on the concepts involved, it follows that the sentence in question encodes information that is knowable only a posteriori.

A related argument against locating information value even only partly in conceptual content is the argument from error. The conceptual content one associates with the name 'Shakespeare' can include varying amounts of misinformation. In extreme cases, one's concept of Shakespeare may be riddled with misattribution and misdescription, enough so as to befit someone else, say Francis Bacon, far better than Shakespeare. Even so, the sentence 'Shakespeare wrote *Timon of Athens*' does not encode misinformation concerning Bacon. If it did, it would be

false. But it is true; it encodes correct information concerning Shakespeare. The conceptual content attached to this sentence may be error ridden. Even so, the information semantically encoded is completely error free.

5.2 *Contextual Theories*

What, then, is information value if not even only partly conceptual content? Alternative proposals have been made. It has often been suggested that the information value of a name, a demonstrative, or a single-word indexical is, broadly put, somehow *contextual* in nature. One such proposal identifies the information value of a name with the linguistic "causal" chain of term acquisition leading from an initial dubbing, in which the name is given to its bearer, to the current use of the name.[8] Another proposal identifies the information value of a term with the criteria for the term's application employed by the linguistic community's "experts" in the use of the term—the experts' conceptual content, as it were.[9] In the case of indexicals (as completed by accompanying demonstrations where necessary), it has been suggested that the term's information value be identified with its character (or, from the vantage point of the modified naive theory, with its program).[10]

Taken individually, each of these theories has its difficulties. The theory that information value is the experts' conceptual content is subject to at least the argument from error urged above, if not to the other objections as well. Moreover, the experts may disagree among themselves. There may even be radical disagreement over the most fundamental aspects of the object so named. There are names of historical figures about whom no one alive is an expert, and there may be names of objects about which no one living or dead knows very much at all. Are there no experts in such cases, or do we all count equally as experts?

An unsatisfactory feature of the proposal to identify the information value of an indexical with its character is that it does not generalize to the original problem case, that of proper names such as 'Hesperus' and 'Phosphorus'. Even if restricted to the special case of indexicals, the theory that information value = character is subject to the twin-earth argument given above; the sentence 'Hubert weighs exactly 165 pounds' is simply replaced with '*He* weighs exactly 165 pounds'.

The theory that the information value of a name is the linguistic network or chain that secures the referent is, on some versions, subject to the argument from subjectivity, since different users of the name enter into different chains of communication. More important, the theory seems ill conceived if not downright desperate. Whereas there is some-

thing natural and compelling about the idea that properties, relations, and the objects that have them and stand in them are the building blocks of information, there is some thing wildly bizarre about the idea that relevant sorts of linguistic chains and causal networks function as building blocks of information. These linguistic chains are typically invoked in the theory of reference to explain how a name, as used by a particular speaker, secures its referent in the user's idiolect. To suppose that these linguistic chains are also information components is a confusion, on the order of a category mistake. The contextual mechanism by which reference in an idiolect is secured is one thing; cognitive information content is quite another.

5.3 *Verbal Theories*

I have argued that the conceptual and contextual theories of information value are unsatisfactory as replacements for the modified naive theory. Even combination theories built from these, such as the theory that information value is partly referent and partly conceptual content, are subject to some of the same objections. In the case of proper names, there is another possibility to be considered: the theory that the information value of a name is, at least in part, simply the name itself. On this theory, the sentence 'Hesperus is Phosphorus' encodes different information than the sentence 'Hesperus is Hesperus' because the information encoded by the former is made up in part of the names 'Hesperus' and 'Phosphorus', and presumably the 'is' of identity, whereas the information encoded by the latter is made up only of the name 'Hesperus' taken twice and the 'is' of identity. Pieces of information are regarded on this view as linguistic objects, made up at least in part of words. Let us call this the *verbal theory of information value*.

The verbal theory of information value fares considerably better on several counts than the conceptual and contextual theories. Unfortunately, though, it too is subject to serious objections. The simplest version of the verbal theory holds that the information value of a name is just the name itself, qua syntactic sound and shape. This simple version is refuted by the twin-earth argument, or, more simply, by the commonplace phenomenon of two individuals having the same name (qua syntactic sound and shape). But the verbal theory of information value need not see information as merely empty syntactic sounds and shapes. A more plausible version of the theory will involve taking the expressions that make up information in a certain way. For example, it might be held that the information value of a name, on a given occasion of use, is the name *with* the referent that it has on that occasion, or, what comes to the same thing, a complex (e.g. an ordered couple)

consisting of the expression qua syntactic sound and shape together with the referent of the expression on that occasion. This rendering of the verbal theory is considerably more plausible, but there are serious objections to it as well. Unless this theory is extended to other sorts of expressions in addition to proper names, we are still left wondering what the information values of the other sorts of expressions (e.g. demonstratives) are. Indeed, the theory cannot be plausibly restricted to proper names, but must be extended at least to those expressions of natural language that are generally regarded as semantically analogous to proper names, such as single-word natural-kind terms. The theory would thus hold that the information encoded by the sentence 'This is water' is made up in part of the word 'water' as used with reference to the substance water, and similarly for sentences involving single-word species terms (e.g. 'tiger') and other single-word natural-kind terms.

The main problem with this theory is that it does not allow for even the possibility of two distinct names or natural-kind terms having the same information value. For example, the information encoded by the English sentence 'Cats have whiskers' cannot on this theory be encoded by a sentence of another language unless that language happens to employ the very expressions 'cat' and 'whisker' just as they are employed in English. This is obviously false. The information that cats have whiskers—the information encoded in English by the sentence 'Cats have whiskers'—is also expressible in any number of languages, regardless of their terms for 'cat' and 'whisker'. Worse yet, it is not in the least bit clear that the verbal theory of information value can be plausibly restricted even to proper names and single-word natural-kind terms. Kripke, Putnam, and others have shown that single-word terms for natural phenomena (such as 'hot') and even single-word non-natural-kind terms (perhaps even artifact terms, such as 'pencil') are also semantically analogous in important respects to proper names. Viewed in this light, the verbal theory of information value may involve the preposterous consequence that no information, or precious little, can be expressed simultaneously in distinct natural languages.[11] In effect, the theory identifies (in at least a wide range of cases) the encoded with the encoder, to the extent that different means of encoding information, if they are syntactically distinct, are held to encode different information. But surely it is central to the intuitive idea of encoding information that any piece of information that can be encoded at all can be encoded systematically by any number of syntactically distinct means.

5.4 *Frege's Strategy Generalized*

Whatever the demerits of these various theories of information value taken individually, there is an even more serious problem faced by most, if not all, of them at once. We have seen that Frege's Puzzle is employed by orthodox theorists as a refutation of the modified naive theory that the information value of a name, a demonstrative, or some other single-world singular term is simply its referent. In effect, Frege and his followers have converted Frege's Puzzle into a strategy for refuting (any modification of) the naive theory. The strategy is to find a pair of terms a and b that share a common referent but are such that $\ulcorner a = b \urcorner$ is informative whereas $\ulcorner a = a \urcorner$ is not. The refutation then proceeds by way of Frege's Law and the compositionality principle mentioned in the preceding chapter. Let us call this *Frege's Strategy*.

What Frege and his followers have failed to notice is that this strategy generalizes. Simply put, the general strategy is this: Let F be any function of a term (e.g. reference) that is identified with information value by the theory to be refuted. Now find a pair of terms a and b sharing the same F, but such that $\ulcorner a = b \urcorner$ is informative whereas $\ulcorner a = a \urcorner$ is not. Then apply Frege's Law and the compositionality principle. The theory that a term's information value is its F is thus refuted. Let us call this the *Generalized Frege Strategy*. It is remarkably general.

If Frege's Strategy works against the modified naive theory, the Generalized Frege Strategy works equally well against a very wide range of theories. It is possible, for example, to apply the Generalized Frege Strategy against some of the contextual theories of information value mentioned above, and against most combination theories of information value. Consider the plausible theory that the information value of a demonstrative, as used on a particular occasion, is partly the associated visual appearance and partly the referent. One sort of example employing the Generalized Frege Strategy is the following. Suppose that you know that Paul, the man standing in front of us, has an identical twin Peter, whom you cannot distinguish from Paul. If I blindfold you for 30 seconds, then release the blindfold, and you see a man looking just like Paul standing in front of us just as Paul was 30 seconds ago, you have no way of knowing with certainty which twin he is. Let us suppose that in fact it is still Paul standing in front of us. If I utter the sentence 'He is him' very slowly, taking the full 30 seconds from start to finish—pointing to Paul while uttering 'He', then blindfolding you, then removing the blindfold, and then pointing to the man in front of us while uttering 'is him'—I will have spoken informatively. But if I do not blindfold you, and I utter the very same sentence, or if instead I simply point to Paul with both hands and utter the same sentence,

my utterance is utterly uninformative, or at least no more informative than 'Hesperus is Hesperus'. In both cases, the visual appearances associated with the demonstratives, as completed by the accompanying ostensions to Paul, are one and the very same, as are the referents. Following the Generalized Frege Strategy, we ought to conclude that the information value of the demonstrative, as thus completed by ostension, is something other than the associated visual appearance, the referent, or even the combination of the two.

Ironically, it may even be possible to turn this form of argument against the original Fregean theory that information value is purely conceptual content. I, like Putnam, do not have the slightest idea what characteristics differentiate beech trees from elm trees, other than the fact that the term for beeches is 'beech' and the term for elms is 'elm'. The conceptual content that I attach to the term 'beech' is the same that I attach to the term 'elm', and it is a pretty meager one at that. My concept of an elm tree is no different from my concept of a beech tree. Nevertheless, it would be news to me to be told that elms and beeches are the very same things. In fact, I know that they are not the same things. At the same time, of course, I know that elms and elms are the same things. Following the Generalized Frege Strategy, we should conclude that the information value of 'elm' or 'beech' is not the conceptual content.[12]

In this application of the Generalized Frege Strategy, the relevant informative identity statement is not even true. The truth of an informative identity statement is required only in the application of Frege's original strategy against theories that locate information value, at least in part, in reference. In the general case, only informativeness is required, and false identity statements are always informative (so informative, in fact, as to be misinformative).

What these and other suitably modified applications of the Generalized Frege Strategy seem to show is that the information values of names, demonstratives, and other single-word indexical singular terms (as well as those of single-word natural-kind terms, and certain other other single words) are more fine-grained than the referents, the purely conceptual contents, the characters, or any combinations thereof. Expressions apparently differing in informative value (assuming Frege's Law and the compositionality principle) may nevertheless share the same referent, the same purely conceptual content, and the same character.

The only candidate for information value suggested so far that might be fine-grained enough is the expression itself. But we have seen that the expressions themselves are *too* fine-grained, since one and the same piece of information may be encoded using syntactically distinct

expressions (e.g. in different languages). However, the Generalized Frege strategy may even be employed against the verbal theory of information value, and this appears to show that even proper names, as used with a particular referent, are not only too fine-grained but also too coarse-grained. Consider the sentence 'Aristotle is identical with Aristotle'. If both occurrences of 'Aristotle' are used with reference to the celebrated philosopher of antiquity, the sentence so understood is uninformative. If the first occurrence of 'Aristotle' is used with reference to the philosopher and the second with reference to the late shipping magnate who married Jacqueline Kennedy, the sentence so understood is informative—indeed preposterous. Suppose for the sake of the example that, though no one now alive knows it, the celebrated philosopher of antiquity did not die in 322 B.C. as we think, but instead went into hiding in Chalcis, discovered the philosopher's stone, which drastically slows down the aging process, and reemerged in the twentieth century as the wealthy and powerful shipping magnate Aristotle Onassis. This would make the sentence understood in the second way true rather than false, but certainly no less informative. It does not affect the uninformativeness of the same sentence understood the first way. Thus, not only may the same information value be attached to syntactically distinct names for Aristotle (e.g. 'Aristotle' and 'Ἀριστοτέλης'), but also, by the Generalized Frege Strategy, it seems that the single name 'Aristotle', as used with reference to the philosopher, could turn out to have distinct information values simultaneously!

We are left, then, with the following uncomfortable situation: If the original argument employing Frege's Strategy is successful against the modified naive theory that the information value of a name is its referent, then structurally analogous arguments employing the Generalized Frege Strategy apply equally successfully against a very wide range of theories of information value, including Frege's own. Virtually any substantive theory of information value imaginable reintroduces a variant of Frege's Puzzle—or else it is untenable on independent grounds (modal or epistemological arguments, the argument from error, the argument from subjectivity, and so on). Just as Frege's Strategy converts Frege's Puzzle into an argument against the modified naive theory of information value, there are formally analogous puzzles that raise equally serious difficulties in connection with nearly any plausible, minimally specific theory of information value that might be dreamt up. This is extremely awkward. One might conclude from this that the information value of a name or indexical is utterly sui generis, with nothing more to be said in the way of illumination other than various negative characterizations to the effect that it is not the referent, not the purely conceptual content, not the character, and so forth. We seem to be able to say what it is

not, but not what it is. Whatever is offered in the way of an illuminating positive characterization may fall prey to the Generalized Frege Strategy. This would be to adopt the defeatist attitude that Frege's challenging questions are in principle unanswerable, or at best, are susceptible only to unilluminating and largely negative characterizations of information value that leave the notion quite obscure.

Suppose we insist that there must be some positive response to Frege's questions, some illuminating account of what the information value of a name is. Since almost any such illuminating account would be refuted by the Generalized Frege Strategy if the naive theory were refuted by Frege's original strategy, we would have to deny that Frege's original strategy actually succeeds in refuting the modified naive theory. Can it be that Frege's Strategy and its generalization are somehow fallacious, and that the information value of a name is its referent after all? It can. In fact, it is.

6

The Crux of Frege's Puzzle

6.1 The Minor Premise

There are three main elements in Frege's Puzzle, and in the corresponding strategy: Frege's Law, the compositionality principle, and the further premise that $\ulcorner a = b \urcorner$ is informative and a posteriori whereas $\ulcorner a = a \urcorner$ is not. I have argued that there is nothing to be gained by challenging the compositionality principle, and that Frege's Law is beyond challenge, since properly understood it is simply a special instance of Leibniz's Law. Still to be considered is the minor premise that $\ulcorner a = b \urcorner$ is informative whereas $\ulcorner a = a \urcorner$ is not.

Historically, philosophers who have had some inclination toward something like the naive theory, including Frege, Mill, and Russell, have allowed that $\ulcorner a = b \urcorner$ is informative and a posteriori whereas $\ulcorner a = a \urcorner$ is not. This was thought too obvious to be denied, and other means for coming to grips with Frege's Puzzle were sought and devised. In contemporary philosophy, direct-reference theorists—who should find the naive theory particularly congenial—have typically conceded this point, or something tantamount to it, and have therefore abstained from outright, unequivocal endorsement of the naive theory or any modification of the naive theory. Consider the following remarks:

> [You] see a star in the evening and it's called 'Hesperus'. . . . We see a star in the morning and call it 'Phosphorus'. Well, then we find . . . that Hesperus and Phosphorus are in fact the same. So we express this by 'Hesperus is Phosphorus'. Here we're certainly not just saying of an object that it's identical with itself. This is something that we discovered. (Saul Kripke, *Naming and Necessity*, pp. 28–29)

> [We] do not know *a priori* that Hesperus is Phosphorus, and are in no position to find out . . . except empirically. (ibid., p. 104; see also the disclaimer on pp. 20–21)

> Before appropriate empirical discoveries were made, men might have failed to know that Hesperus was Phosphorus, or even to believe it, even though they of course knew and believed that Hesperus was Hesperus. (Kripke, "A Puzzle About Belief," p. 243—but see p. 281, note 44; see also the disclaimer at p. 273, note 10)

> Certainly Frege's argument shows meaning cannot just *be* reference. . . . (Hilary Putnam, "Comments," p. 285)

> If we distinguish a sentence from the proposition it expresses then the terms 'truth' and 'necessity' apply to the proposition expressed by a sentence, while the terms '*a priori*' and '*a posteriori*' are sentence relative. Given that it is true that Cicero is Tully (and whatever we need about what the relevant sentences express) 'Cicero is Cicero' and 'Cicero is Tully' express the same *proposition*. And the proposition is necessarily true. But looking at the proposition through the lens of the *sentence* 'Cicero is Cicero' the proposition can be seen *a priori* to be true, but through 'Cicero is Tully' one may need an *a posteriori* investigation. (Keith Donnellan, "Kripke and Putnam on Natural Kind Terms," note 2 on p. 88)

> Faced with Frege's identity puzzle, it is difficult indeed to maintain that the names 'Hesperus' and 'Phosphorus' make precisely the same contribution to the information content of sentences that contain either one. Such a claim would be extremist. (Nathan Salmon, *Reference and Essence*, p. 13)

Here is where well-intentioned philosophers have been led astray. It is precisely the seemingly trivial premise that $\ulcorner \bar{a} = \bar{b} \urcorner$ is informative whereas $\ulcorner \bar{a} = \bar{a} \urcorner$ is not informative that should be challenged, and a proper appreciation for the distinction between semantically encoded and pragmatically imparted information points the way. Recall that Frege's Law is erected into a truth of logic by understanding the word 'informative' in such a way that to say that a sentence is informative is to say something about its information content. By the same token, however, with 'informative' so understood, and with a sharp distinction between semantically encoded information and pragmatically imparted information kept in mind, it is not in the least bit obvious, as Frege's Puzzle maintains, that $\ulcorner \bar{a} = \bar{b} \urcorner$ is, whereas $\ulcorner \bar{a} = \bar{a} \urcorner$ is not, informative *in the relevant sense*. To be sure, $\ulcorner \bar{a} = \bar{b} \urcorner$ *sounds* informative, whereas $\ulcorner \bar{a} = \bar{a} \urcorner$ does not. Indeed, an utterance of $\ulcorner \bar{a} = \bar{b} \urcorner$ genuinely imparts information that is more valuable than that imparted by an utterance of $\ulcorner \bar{a} = \bar{a} \urcorner$. For example, it imparts the nontrivial linguistic information

about the sentence $\ulcorner a = b \urcorner$ that it is true, and hence that the names *a* and *b* are co-referential. But that is pragmatically imparted information, and presumably not semantically encoded information. (See the discussion in section 3.2 of the "Begriffsschrift" solution to Frege's Puzzle.) It is by no means clear that the sentence $\ulcorner a = b \urcorner$, stripped naked of its pragmatic impartations and with only its properly semantic information content left, is any more informative in the relevant sense than $\ulcorner a = a \urcorner$. Abstracting from their markedly different pragmatic impartations, one can see that these two sentences may well semantically encode the very same piece of information. I believe that they do. At the very least, it is by no means certain, as Frege's Puzzle pretends, that the difference in "cognitive significance" we seem to hear is not due entirely to a difference in pragmatically imparted information. Yet, until we can be certain of this, Frege's law cannot be applied and Frege's Puzzle does not get off the ground. In effect, then, Frege's Strategy begs the question against the modified naive theory. Of course, if one fails to draw the distinction between semantically encoded and pragmatically imparted information, as so many philosophers have, it is small wonder that information pragmatically imparted by (utterances of) $\ulcorner a = b \urcorner$ may be mistaken for semantically encoded information.[1] If Frege's Stategy is ultimately to succeed, a further argument must be made to show that the information imparted by $\ulcorner a = b \urcorner$ that makes it sound informative is, in fact, semantically encoded. In the meantime, Frege's Puzzle by itself is certainly not the final and conclusive refutation of the modified naive theory that the orthodox theorists have taken it to be. For all that Frege's Strategy achieves, the modified naive theory remains the best and most plausible theory available concerning the nature and structure of the information encoded by declarative sentences.

Ironically, as was noted in section 4.2, Frege was not unaware of the distinction between semantically encoded and merely pragmatically imparted information. He did not fully appreciate the significance of this distinction for his theory of information content. In particular, he failed to notice that the distinction undermines his main argument against the naive theory.

6.2 *Substitutivity*

The general puzzle, however, is not so easily put to rest. Although the premise that $\ulcorner a = b \urcorner$ is informative whereas $\ulcorner a = a \urcorner$ is not facilitates the derivation of Frege's Puzzle, this premise is not an essential element in the general puzzle. The premise is invoked in conjunction with Frege's Law to establish the result that there are pairs of sentences of

the form ϕ_a and ϕ_b that differ in information content from one another—i.e., that encode different pieces of information—even though a and b are co-referential (genuine) proper names, demonstratives, single-word indexical singular terms, or any combination thereof. This is the crux of Frege's Puzzle. One might attempt to establish this result in some more general way, without invoking the suspect premise that $\ulcorner a = b \urcorner$ is informative. As Michael Dummett has stressed, and as Frege's formulation of the puzzle clearly indicates, the notion of *information content* relevant to Frege's Puzzle is closely tied to the ordinary, everyday notions of *knowledge* and *belief*. One intuitively appealing picture that is entrenched in philosophical tradition depicts belief as a type of inward assent, or a disposition toward inward assent, to a piece of information. To believe that p is to concur covertly with, to endorse mentally, to nod approval to, the information that p when p occurs to you. At the very least, to believe that p one must adopt some sort of favorable disposition or attitude toward the information that p. In fact, the adoption of some such favorable attitude toward a piece of information is both necessary and sufficient for belief. That is just what belief is.[2] To believe that p is, so to speak, to include that piece of information in one's personal inner "data bank." It is to have that information at one's disposal to rely upon, to act upon, to draw inferences from, or to do nothing with. Belief is thus a relation to pieces of information.

These observations suggest the following principal schema, where the substituends for S and S' are declarative English sentences:

> If the information that S = the information that S', then someone believes that S if and only if he or she believes that S'.

Analogous schemata may be written for assertion and the other so-called propositional attitudes of knowledge, hope, and so forth. Like Frege's Law, each of these schemata may be regarded as (formal mode renderings of) so many instances of Leibniz's Law. In fact, Frege's Law can be viewed as a minor variation of one such schema:

> If the information that S = the information that S', then it is informative (knowable only a posteriori, a valuable extension of our knowledge, etc.) that S if and only if it is informative (a posteriori, etc.) that S'.

The *thesis of the substitutivity of co-informational sentences in propositional attitude contexts* is the thesis that every proper instance of any of these schemata is true. This may be separated into the *thesis of the substitutivity of co-informational sentences in assertion contexts* and so on for each of the attitudes. The thesis, or theses, is virtually a logical consequence of the idea that the object or content of a given belief,

piece of knowledge, etc., is a piece of information, or a "proposition", and that a sentence encoding that information thereby gives the content of the belief. This idea, or something like it, is a commonplace in the philosophy of language; it is usually taken for granted without challenge by both sides in philosophical disputes over related issues (such as the question of the logical form of belief attributions). Some philosophers, in an effort to rescue a favored theory of propositions from the pitfalls of propositional attitude contexts, have rejected the thesis of substitutivity of co-informational (or co-propositional) sentences in propositional attitude contexts. But doing so seems both extreme and ad hoc. If the favored theory of propositions conflicts with the thesis, it would be more plausible to reject the theory.[3]

Insofar as some of the substitutivity theses are accepted as plausible principles concerning the relation between the pieces of information contained in a sentence and the content of an attitude (belief, knowledge, etc.) thereby expressed, they yield an important procedure for establishing that two given pieces of information are distinct. One may simply rely on our ordinary, everyday criteria, whatever they happen to be, for correctly saying that someone believes or knows something or does not believe or know it. We do not have to be able to specify these criteria; we need only to be able to apply them correctly in certain paradigm cases.

Now, there is no denying that, given the proper circumstances, we say things like 'Lois Lane does not realize (know, believe) that Clark Kent is Superman' and 'There was a time when it was not known that Hesperus is Phosphorus'. Such pronouncements are in clear violation of the modified naive theory taken together with the thesis of substitutivity of co-informational sentences in doxastic and epistemic contexts. When we make these utterances, we typically do not intend to be speaking elliptically or figuratively; we take ourselves to be speaking literally and truthfully. Of course, one could intentionally utter such sentences in a metaphorical vein, or as an ellipsis for something else, but such circumstances are quite different from the usual circumstances in which such utterances are made, which are so familiar to teachers and students of contemporary analytic philosophy. The crucial question, however, is whether when we say such things we are correctly applying the criteria that govern the correct use of propositional-attitude locutions.

Recently a number of philosophers, mostly under the influence of the direct reference theory, have expressed doubt about the literal truth of such utterances in ordinary usage. If someone believes that Hesperus is a planet, they claim, then, strictly speaking, he or she also believes that Phosphorus is a planet, regardless of what the philosophically untutored or unenlightened say about his or her belief state. Whatever

fact such speakers are attempting to convey by denying the belief ascription, the fact is not the lack of the ascribed belief but something else—perhaps the lack of a corresponding metalinguistic belief to the effect that a certain sentence is true. It is my view that this general approach to these problems is essentially correct, as far as it has been developed. The major problem with this approach is that it has not been developed far enough. I shall say more about this in due course. First, however, it is important to note a glaring philosophical difficulty inherent in this approach.

It is easy nowadays to get caught up in direct-reference mania, but one should never be blinded to possible departures from standard and generally reliable philosophical method and practice. What is ordinarily said in everyday language about a certain set of circumstances—where we take ourselves to be speaking literally and truthfully, and where the circumstances are judged to constitute a paradigm case of what we are saying, etc.—is often regarded as an important datum, sometimes the only possible datum, relevant to a certain philosophical or conceptual question about the facts in the matter. Of course, what we ordinarily say in everyday language is sometimes misleading, sometimes irrelevant, sometimes just plain wrong, but in cases where the issue concerns the applicability or inapplicability of a certain concept or term ordinary usage is often the best available guide to the facts. Consider, for example, the sorts of considerations invoked by epistemologists in deciding that Edmund Gettier's celebrated examples constitute genuine counterexamples to the traditional analysis of knowledge as justified true belief, or the sorts of considerations invoked by philosophers of perception in deciding that the state of experiencing a visual impression that is in fact caused by and resembles a certain external object is not the same thing as *seeing* the object. In the familiar problem cases, we simply do not say that the subject *knows* the relevant piece of information, or that he or she *sees* the relevant object. That is not the way we speak. Our forbearance in attributing knowledge or visual perception in these cases is rightly taken as conclusive evidence that such attributions are strictly false, given the actual and ordinary meanings of 'know' and 'see'. Philosophical programs such as that of analyzing knowledge or that of analyzing perception are, in a significant sense, at least partly an attempt to specify and articulate the implicit criteria or principles that govern the correct application of such terms as 'know' and 'see'. It is precisely for this reason that philosophers so often consult linguistic intuition in doing epistemology or metaphysics. Ordinary language is relevant because it is, at least to some extent, ordinary language that is under investigation. And ordinary usage is a reliable guide to the principles governing the correct use of ordinary language. When the

traditional analyses of knowledge or perception are challenged through thought experiments concerning what we would say in certain problem cases, philosophers are rightly skeptical of the reply that ordinary usage is incorrect and that the subject does indeed know the proposition in question, or see the object in question, even though we typically say that he or she does not. Anyone maintaining this position may well by suspected of protecting an invested interest in the theory being challenged, rather than pursuing in good faith the philosopher's primary purpose of seeking truth no matter where the facts may lead. This is not to disparage such concepts as *justified true belief* and *experiencing a visual impression caused by and resembling an external object*. Such concepts may be epistemologically important. However, they demonstrably do not correspond—at least, they do not correspond exactly—to the everyday criteria that are implicit in ordinary usage for knowing or seeing. These criteria are, in a significant sense, *what are in question*.

Similarly, the claim that Lois Lane does, strictly speaking, believe and even know that Clark Kent is Superman (since she knows that he is Clark Kent) must not be made lightly, lest he or she who makes it be placed under the same suspicion. For here the question concerns, at least partly, the tacit principles governing the correct use of ordinary-language words such as 'believe', and the ordinary-usage evidence against the claim is strong indeed. The plain fact is that we simply do not speak that way. Perhaps we should learn to use a language in which propositional-attitude idioms function in strict accordance with the modified naive theory across the board, including the troublesome 'Hesperus'-'Phosphorus' and 'Cicero'-'Tully' cases, since ordinary language already agrees with the modified naive theory in the other, more commonplace sorts of cases. But that is a question for prescriptive philosophy of language, not one for descriptive philosophy of language. The more immediate and pressing philosophical question concerns the actual criteria that are implicitly at work in the everyday notion of belief, and the other attitudes, in their crude form, as they arise in real life without theoretical or aesthetic alteration.

I maintain that, according to these very criteria (in the standard sort of circumstance), it is, strictly speaking, correct to say that Lois Lane does know that Clark Kent is Superman, and that when ordinary speakers deny this they are typically operating under a linguistic confusion, systematically misapplying the criteria that govern the applicability or inapplicability of their own doxastic and epistemic terms and concepts. Similarly, anyone who knows that Hesperus is Hesperus knows that Hesperus is Phosphorus, no matter how strongly he or she may deny the latter. Moreover, anyone who knows that he or she knows that

Hesperus is Hesperus also knows that he or she knows that Hesperus is Phosphorus, no matter how self-consciously he or she may disbelieve that Hesperus is Phosphorus.[4]

These claims clash sharply with ordinary usage. Whereas it is (as I have argued) extremely important not to lose sight of the tried and true philosophical tool of looking to ordinary usage in such matters, it is equally important to recognize the limitations of that test. Ordinary usage is a reliable guide to correct usage, but it is only a guide. Ordinary usage can sometimes be incorrect usage. Even when the ordinary usage of a certain locution is systematic, it can be systematically incorrect—if, for example, the language is deficient in ways that compel speakers to violate its rules in order to convey what they intend, or if the principles and social conventions governing the appropriateness of certain utterances require certain systematic violations of the principles and rules governing correct and incorrect applications of the terms used. My claim is that ordinary usage with regard to such predicates as 'is aware that Clark Kent is Superman' and 'believes that Hesperus is Phosphorus' conflicts with the criteria governing their correct application in just this way. However inappropriate it may be in most contexts to say so, Lois Lane is (according to the myth) fully aware that Clark Kent is Superman, and anyone who believes that Hesperus is Hesperus does in fact believe that Hesperus is Phosphorus. We do not speak this way; in fact, it is customary to say just the opposite. But if we wish to utter what is true, and if we care nothing about social convention, we should speak this way. The customary way of speaking involves us in uttering falsehoods.

Of course, it is no defense of the modified naive theory simply to make these bold claims. It is incumbent on the philosopher who makes these claims (i.e., me) to offer some reason for supposing that ordinary speakers, in the normal course of things, would be led to distort the rules of language systematically, so that ordinary usage cannot be relied upon in these cases as a guide to the correct-applicability conditions of the relevant terms and concepts. The account I shall offer is complex. The main part of this account will be given in section 8.4. For now, a tentative account is provided by repeating the distinction between semantically encoded and pragmatically imparted information. If one is not careful to keep this distinction in mind, it is altogether too easy to confuse information pragmatically imparted by (utterances of) 'Hesperus is Phosphorus' for semantically encoded information. In saying that *A* believes that Hesperus is Phosphorus, taken literally, we are merely attributing to *A* a relation (belief) to a certain piece of information (the information semantically encoded by 'Hesperus is Phosphorus'). The 'that'-clause 'that Hesperus is Phosphorus' functions here as a means for referring to that piece of information. Since the form of words

'Hesperus is Phosphorus' is considerably richer in pragmatic impartations than other expressions having the same semantic information content (e.g. 'Hesperus is Hesperus'), if one is not careful one cannot help but mistake the 'that'-clause as referring to this somewhat richer information—information which *A* may not believe. (See note 1.) Utterances of the locution $\ulcorner a$ believes that $S\urcorner$ may even typically involve a Gricean implicature to the effect that the person referred to by *a* believes the information that is typically pragmatically imparted by utterances of *S*. Even so, that is not part of the literal content of the belief attribution. The general masses, and most philosophers, are not sufficiently aware of the effect that an implicature of this kind would have on ordinary usage. It is no embarrassment to the modified naive theory that ordinary speakers typically deny literally true belief attributions (and other propositional-attitude attributions) when these attributions involve a 'that'-clause whose utterance typically pragmatically imparts information which the speaker recognizes not to be among the beliefs (or other propositional attitudes) of the subject of the attribution. In fact, it would be an embarrassment to the modified naive theory if speakers did not do this. With widespread ignorance of the significance of the distinction between semantically encoded and pragmatically imparted information, such violation of the rules of the language is entirely to be expected.

7

More Puzzles

7.1 The New Frege Puzzle

The distinction between semantically encoded and pragmatically imparted information goes a long way toward solving the problems posed by Frege's Puzzle and the apparent failure of substitutivity of proper names and other single-word singular terms in propositional-attitude contexts. There can be little doubt that failure to appreciate the distinction is largely responsible for the relative unpopularity of the modified naive theory in favor of its rivals throughout the history of the theory of meaning. Unfortunately, the distinction does not yield the final word on the general problem. A version of this general problem arises again, this time in a particularly strengthened form, when one takes note of the following fact: Even a speaker who has been fully apprised of the distinction between semantically encoded and pragmatically imparted information, and who has learned to be scrupulously careful about separating out pragmatic impartations when dealing with matters of semantics, may give assent to some sentence *S* which encodes a certain piece of information and which the speaker fully understands, while the same speaker may fail to give assent, and may even give dissent, to some sentence *S'* which the speaker also fully understands and which, according to the modified naive theory, encodes the very same information. This can easily happen even if the speaker is perfectly rational, mentally acute (in fact, an ideally perfect thinker), eager to indicate his or her beliefs through verbal assent and dissent, and a firm and dogmatic believer in the modified naive theory!

In saying that someone fully understands a sentence, I mean only that he or she associates the right proposition with the sentence in the right way (that is, unconsciously "computes" the semantically encoded content of the sentence from the recursion rules of semantic composition, or something along these lines—however it is that we get things right when we understand a sentence), and that he or she has a complete grasp of this proposition. In particular, knowing the truth value of the proposition is not required for complete understanding.

For example, suppose that Lois Lane is forced to endure a full academic year of intensive training in the theory of meaning through the writings of a famous Kryptonian philosopher of language. On Krypton (Superman's native planet, according to the myth), the distinction between semantically encoded and pragmatically imparted information was duly appreciated, and the modified naive theory was held in the highest esteem by all but a very small minority of semanticists. The modified naive theory is drilled into her head. She is instructed in the distinction between semantically encoded and pragmatically imparted information, and she is taught to assent to all and only those sentences whose semantically encoded information content she believes and to dissent from all and only those sentences whose negation commands her assent. Now consider the following two sentences:

(5) Superman fights a never-ending battle for truth, justice, and the American way.

(6) Clark Kent fights a never-ending battle for truth, justice, and the American way.

If anyone understands these sentences, Lois does. She fully grasps the proposition encoded by these sentences, and she associates the right proposition with each sentence. One might wonder whether she fully understands sentence 6, but a moment's reflection confirms that she does. For example, she certainly does not misunderstand sentence 6 to mean that Perry White is a tyrant. She correctly understands sentence 6 to mean that Clark Kent fights a never-ending battle for truth, justice, and the American way. Lois grasps this information as well as anyone does. Of course, she wrongly believes it to be misinformation, but getting clearer about its truth value would not enable her to grasp it any deeper. So Lois correctly understands both sentences. Yet she verbally assents to sentence 5 and verbally dissents from sentence 6. The fact that she fails to assent to, and in fact dissents from, sentence 6 when she correctly understands it to mean that Clark Kent fights a never-ending battle for truth, justice, and the American way, is very strong evidence that she does not believe this information. This is especially true if one takes seriously the analysis of belief suggested in the preceding section, whereby belief is identified with inward assent or agreement to a piece of information or with a disposition toward inward assent. Given that Lois sincerely wishes to reveal her opinions through verbal assent and dissent, that she correctly understands what is meant by sentence 6, and that she is a perfectly rational and competent thinker, her verbal dissent from sentence 6 would seem to be as good an indication as one could possibly have that she inwardly dissents

from the proposition. If she inwardly assented to the proposition, it would seem, she would outwardly assent to the sentence. Her failure to assent to sentence 6, therefore, provides an extremely compelling reason to suppose that she does not believe what she correctly understands it to mean. Similarly, Lois's assent to sentence 5 provides extremely compelling evidence, evidence as good as one could ever have, that she believes this piece of information. Her combined verbal behavior, then, provides an extremely compelling reason to conclude that she believes that Superman fights a never-ending battle for truth, justice, and the American way, but does not believe that Clark Kent does. No doubt, this is also part of the original justification for saying just this about Lois's beliefs. This characterization of Lois's beliefs flatly contradicts the modified naive theory.

It is no help to appeal here to ignorance of the distinction between semantically encoded and pragmatically imparted information, for both Lois (whose beliefs we are talking about) and we (who are talking about those beliefs) are by now well aware of the distinction. Awareness of the distinction does nothing to obviate the compelling force of the evidence provided by Lois's verbal behavior. In particular, it does nothing to dissipate the extremely compelling grounds, provided by Lois's failure to assent to sentence 6, for concluding that she does not believe that Clark Kent fights a never-ending battle.

These considerations generate another puzzle for the modified naive theory. It was argued that Lois correctly, completely, and fully understands both sentence 5 and sentence 6. In particular, she correctly understands sentence 6 to mean that Clark Kent fights a never-ending battle for truth, justice, and the American way. Which proposition does she take sentence 6 to encode? Given her working knowledge of English, her acquaintance with Clark Kent, and her recent training in the philosophy of language, it can only be the singular proposition about Clark Kent that he fights a never-ending battle for truth, justice, and the American way. Now, according to the modified naive theory, Lois believes this singular proposition, for she believes of Superman that he fights a never-ending battle for truth, justice, and the American way. If anyone is ever in a position to have de re beliefs about Superman, Lois has this particular de re belief about him. On the modified naive theory, the content of this de re belief simply is the very proposition that she correctly takes sentence 6 to encode. Hence, on the modified naive theory, Lois—whom we may suppose to be an ideally rational and competent speaker and who sincerely wishes to reveal her opinions through verbal assent and dissent—correctly identifies which proposition is encoded by sentence 6, and she firmly believes this very proposition. Yet, even on reflection, she fails to assent to sentence 6, and

in fact dissents from it. What, on the modified naive theory, can account for her behavior? How can the theory explain away her failure to assent to sentence 6 as grounds for concluding that she does not believe that Clark Kent fights a never-ending battle for truth, justice, and the American way?

Let us take a more familiar example. An ancient astronomer-philosopher, well versed in the modified naive theory and the distinction between semantically encoded and pragmatically imparted information, verbally assents to (his sentence for) the sentence 'Hesperus is Hesperus' without assenting to the sentence 'Hesperus is Phosphorus'. It is not enough to explain this phenomenon by pointing out that the astronomer-philosopher does not realize that the second sentence encodes information that he believes, or that the two sentences encode the same information, or that one sentence is true and commands his assent if and only if the other one is and does. The question is: How can he fail to realize any of this? We may suppose (1) that he fully grasps the proposition about the planet Venus and the planet Venus that the former is the latter, (2) that, being an adherent of the modified naive theory, he takes the first sentence to encode this very proposition and no other, and (3) that it is this very same proposition and no other that he also takes the second sentence to encode (since this is also the proposition about Hesperus and Phosphorus that they are identical). How then can he fail to see that the sentences are informationally equivalent? Morover, he fully endorses this proposition, so how, upon reflection, can he fail to be moved to assent to the second sentence when it is this very proposition—one he fully grasps and believes—that he takes the second sentence to encode? The situation becomes especially puzzling for the adherent of the modified naive theory if we suppose that, in believing the proposition that Hesperus is Hesperus, the ancient astronomer-philosopher inwardly assents to it, or is so disposed. If he assents inwardly to the proposition, or is so disposed, why, if he is reflective and eager to reveal his beliefs through verbal assent, is he not similarly disposed to assent outwardly to a sentence which he takes to encode that very proposition? The distinction between semantically encoded and pragmatically imparted information sheds no light on this new problem, for we are supposing that the ancient astronomer-philosopher is well aware of the distinction and never allows himself to be misled by pragmatic impartations in matters concerning semantic content. Moreover, we may also suppose that there is nothing whatsoever wrong or imperfect about the astronomer-philosopher's reasoning or thought processes. We may even suppose him to have superhuman intelligence (or as much intelligence as is compatible with his not knowing the truth of 'Hesperus is Phosphorus'). What, then,

is preventing him from making the connection between what he takes the sentence to encode and his belief of that very information?

It appears that the modified naive theory turns against itself in discourse involving propositions about singular propositions, for, on the modified naive theory, these too are singular propositions. (See chapter 6, note 4.) If the ancient astronomer-philosopher believes that 'Hesperus is Phosphorus' encodes the information that Hesperus is Phosphorus, then, on the modified naive theory, he also believes that 'Hesperus is Phosphorus' encodes the information that Hesperus is Hesperus—information which he fully grasps and firmly believes on logical grounds alone. It seems to follow that the mere understanding of the sentence should suffice to elicit his unhesitating and unequivocal assent, even if he is not so intelligent. But, as Frege rightly noted, there was a time when the mere understanding of this sentence was not sufficient to elicit the assent of astronomers who understood it, and may even have elicited emphatic dissent. This is not a particularly bizarre state of affairs: it is perfectly reasonable that this would be their reaction given the state of ignorance at the time. Yet the modified naive theory seems to lack the means to give a coherent account of this state of affairs without making it appear quite paradoxical.

What we have here is a new and stronger version of Frege's Puzzle, one that does not rely on the question-begging premise that 'Hesperus is Phosphorus' is (semantically) informative, or that someone may believe that Hesperus is Hesperus without believing that Hesperus is Phosphorus, or indeed any premise involving notions such as informativeness or a priority. The new version of the puzzle makes do instead with a weaker, less philosophical-theory-laden, and clearly undeniable premise. The new premise is this:

> Someone who is reflective, without mental defect, and eager to reveal his or her beliefs through verbal assent may correctly identify the information encoded by 'Hesperus is Hesperus', fully grasp that information, indicate concurrence with that information by readily assenting to the sentence, correctly identify the information encoded by 'Hesperus is Phosphorus', fully grasp that information, and yet not feel the slightest impulse to assent to the latter sentence.

In addition, Frege's Law is replaced by the following analogue:

> If a declarative sentence *S* has the very same cognitive information content (Erkenntniswerte) as a declarative sentence *S'*, then an ideally competent speaker who fully understands both sentences perfectly, reflects on the matter, is without mental defect, is eager to indicate his or her beliefs through sincere verbal assent and

dissent, and has no countervailing motives or desires that might prevent him or her from being disposed to assent verbally to a sentence while recognizing its information content as something believed, is disposed to assent verbally to *S* if and only if he or she is disposed to assent verbally to *S'*.

Given, further, the compositionality principle for pieces of information, we have all of the makings of a new and more powerful refutation of the modified naive theory. The distinction between semantically encoded and pragmatically imparted information simply has no bearing on this new argument.

7.2 Elmer's Befuddlement

7.2.1 The Example

The new version of Frege's Puzzle derives its additional strength by invoking dispositions to verbal assent in place of informativeness. We can construct a variant of this stronger version of the puzzle directly in terms of belief without invoking dispositions to verbal assent to sentences. One such variant of the new Frege Puzzle is, in some respects, even stronger than the new Frege Puzzle itself, though ironically it also helps to bring out the modified naive theory's means for solving the general problem. This is best demonstrated by means of a paradox generated by an elaborate example, which I shall call *Elmer's Befuddlement*. Rather than present the entire example all at once, it is more instructive to consider a major part of the example first, in order to test our intuitions about this aspect of the example before considering the example in its entirety.

Elmer's Befuddlement (Excerpts)

> Elmer, a bounty hunter, is determined to apprehend Bugsy Wabbit, a notorious jewel thief who has so far eluded the long arm of the law. Before setting out after Bugsy, Elmer spends several months scrutinizing the FBI's files on Bugsy, studying numerous photographs, movies, and slides, listening carefully to tape recordings of Bugsy's voice, interviewing people who know him intimately, and so on. After learning as much about Bugsy as he can, on January 1 Elmer forms the opinion that Bugsy is (is now, has always been, and will always be throughout his lifetime) dangerous. . . .
>
> On June 1, Elmer receives further information from the FBI that Bugsy was last seen in a club in uptown Manhattan, walking away from a poker game after a gangster type had accused him of cheat-

> ing. This further information gives Elmer pause. He thinks to himself: "Maybe Bugsy . . . is harmless after all. I used to believe that he is a dangerous man, but now . . . I don't know what to think. Maybe he's dangerous, maybe not. I'll just have to wait and see."

Here now is a little two-part quiz: (A) Before June 1, did Elmer believe that Bugsy Wabbit is dangerous? (B) If so, does he continue to believe this even after taking into account the further information he received from the FBI on June 1?

Clearly, question A must be answered affirmatively; Elmer believed for a full five months, from January 1 to June 1, that Bugsy is dangerous, right up until he received the further information concerning Bugsy. This must be so on any reasonable theory of the nature of belief, and it is so on the modified naive theory in particular. On the modified naive theory, to believe that Bugsy is dangerous is to believe the singular proposition about Bugsy that he is dangerous, which is the same thing as believing of Bugsy that he is dangerous. Surely, Elmer had this belief about Bugsy before June 1. If anyone can ever be in a position to have beliefs about Bugsy Wabbit without actually meeting him face to face, then surely Elmer was in such a position when he first decided on January 1 that Bugsy is dangerous. He knew as much about Bugsy as anyone did, save perhaps Bugsy himself, and he may even have known a few things about Bugsy that Bugsy himself did not know.

It would appear equally obvious that question B should be answered negatively. Once he takes the new information into account, Elmer suspends judgment about whether Bugsy is dangerous. Hence, he no longer believes that Bugsy is dangerous. If anyone can ever give up a formerly held belief about someone, Elmer's situation on June 1 would appear to be a typical and central case of such an occurrence. This is not to say, of course, that Elmer now believes that Bugsy is not dangerous, for he does not. Elmer has reconsidered the question of whether Bugsy is dangerous, and he now withholds belief as well as disbelief. Having reconsidered the question, he now believes neither that Bugsy is dangerous nor that he is not. That is what it means to say that Elmer now suspends judgment.

But things are not as clear as they seem. Let us turn now to the example in its entirety.

Elmer's Befuddlement (Unabridged)

> As already recounted, Elmer the bounty hunter forms the opinion on January 1 that Bugsy is (is now, has always been, and will always be throughout his lifetime) dangerous.

Shortly thereafter, having learned that there is a bounty hunter after him, Bugsy undergoes extensive plastic surgery, so that he looks nothing like his former photographs. He also has his voice surgically altered, adopts an entirely new set of personality traits and mannerisms, and so on. He retains his name, however, since it is such a common name.

Hot on Bugsy's tail, Elmer eventually meets up with the new Bugsy Wabbit. Noting that this man is nothing like the Bugsy Wabbit he is pursuing, Elmer falls for Bugsy's ruse and concludes that this Bugsy Wabbit is simply another person with the same name. Elmer befriends Bugsy, but never learns his true identity.

On April 1, Elmer happens to overhear a dispute (apparently over 24 carrots) between Bugsy and someone, and notices that the other party in the dispute is extremely deferential, almost as if he were positively frightened of Bugsy. Elmer decides then and there that this Bugsy Wabbit is also a dangerous man. He says to himself: "I'd better watch my step with my new friend, for Bugsy is a dangerous fellow. In this one respect, the two Bugsy Wabbits are alike."

On June 1, as already recounted, Elmer receives from the FBI further information that gives him pause. He thinks to himself: "Maybe Bugsy the criminal is harmless after all. I used to believe that he is a dangerous man, but now I'm not so sure. In every other respect he is nothing like my friend Bugsy Wabbit, so perhaps I was a bit hasty in deciding that the two Bugsies are like each other in this one respect. My friend Bugsy is definitely dangerous, I haven't changed my mind about that. But as for the jewel thief, I don't know what to think. Maybe he's dangerous, maybe not. I'll just have to wait and see."

Elmer waits, but he never sees. Even today, Elmer feels certain that his friend Bugsy is dangerous, but still wonders whether Bugsy the criminal is dangerous or not.

The saga of Elmer's pursuit of Bugsy Wabbit presents many of the familiar problems. It is reminiscent of Quine's famous example about Ralph and Bernard J. Ortcutt, as well as Kripke's example about Pierre and London, and it has significant points in common with a number of other examples, including Castañeda's examples concerning belief about oneself. There are special aspects of Elmer's Befuddlement that are not present in these other examples, and I shall focus on these special features to construct a paradox.[1]

7.2.2 The Puzzle

Consider again question B: Once Elmer takes account of the further information obtained from the FBI on June 1, how does he stand with respect to the information (or misinformation, as the case may be) that Bugsy Wabbit is dangerous? Does he or does he not believe this piece of information concerning Bugsy?

Let us first consider a simpler question. Roll back the time to April, before Elmer came to have second thoughts about the criminal. Did he believe, at that time, that Bugsy Wabbit is dangerous?

The answer must be that he did. The reasoning that this must be the answer goes as follows: In considering question A, we had already decided that Elmer believed on January 1 that Bugsy is dangerous. We did not yet have the whole story concerning Elmer and Bugsy, but all of the additional information that we have been given concerns events that take place some time after January 1. Hence, the original grounds for claiming that Elmer believes on January 1 that Bugsy is dangerous still obtain. On the modified naive theory in particular, it is still true that Elmer's having familiarized himself with Bugsy's history and appearance in the way he did places him in a position on January 1 to be able to believe at that time of Bugsy that he is dangerous. Now, on April 1 Elmer formed the opinion that his friend Bugsy is dangerous. In doing so Elmer was ignorant of certain critical information concerning Bugsy, but that does not alter the fact that he also steadfastly maintained his view, which he had held since January 1, that Bugsy is dangerous. He did not yet change his mind about Bugsy, first believing him to be dangerous and then giving up that belief. If he believed it before, he believes it still.

There is, it must be admitted, something quite peculiar about Elmer's doxastic state on April 1. There is some sense in which Elmer *comes to* believe on April 1 that Bugsy Wabbit is dangerous (comes to believe of Bugsy that he is dangerous), but there is also some sense in which Elmer believed this about Bugsy since January and never stopped believing it. To give some account of how it is that someone can come to believe something that he or she already believes without ever having ceased to believe it is already a problem for the modified naive theory. I shall not pause here to discuss this. The problem I shall discuss is a sharpened version of this problem, and its solution entails a solution to the present problem. What matters so far is that, however peculiar his doxastic state on April 1, Elmer believed at that time that Bugsy is dangerous.

Now, what about the following summer? Does Elmer continue to believe that Bugsy Wabbit is dangerous even after taking account of the further information from the FBI?

Here no simple 'yes' of 'no' answer by itself is entirely satisfactory. In particular, no simple 'yes' or 'no' answer is satisfactory even if we presuppose the correctness of the modified naive theory. On the one hand, it is critical to the story that in some sense Elmer came to believe on January 1 that Bugsy is dangerous but that Elmer now suspends judgment. Hence, there is an important sense, critical to the story, in which Elmer now believes neither that Bugsy is dangerous nor that he is not dangerous. But it is surely not enough to say that Elmer believes neither that Bugsy is dangerous nor that he is not, and to leave the matter at that, for there is also a very compelling reason to say that Elmer still believes that Bugsy is dangerous: Something exactly analogous to the grounds for holding that Elmer continues to believe on April 1 that Bugsy is dangerous also obtains on June 1. Elmer has not relinquished his opinion that his friend Bugsy is dangerous. If he believed it on April 1, it would seem, he believes it still.

If Elmer had decided on January 1 that Bugsy is dangerous, and had come to have second thoughts on June 1 as he actually did, but had never met Bugsy in the interim and had never formed any further opinion about him, then we would not hesitate to say that Elmer believed on January 1 that Bugsy is dangerous but believes it no longer. Indeed, that is precisely what we did say when we first considered question B, before we knew about Elmer's encounters with Bugsy after January 1. All the information we had given seemed enough to determine that the answer to question B is that Elmer no longer believes that Bugsy is dangerous. Our being given further information concerning Elmer and Bugsy cannot alter what is already determined by the information on hand. If part of the story of Elmer's befuddlement entails that Elmer no longer believes that Bugsy is dangerous, then so does the whole story. (If S entails T, then so does $\ulcorner\overline{S}$ and $\overline{S'}\urcorner$.)

In fact, if Elmer had decided on January 1 that Bugsy is dangerous, and had come to have second thoughts on June 1 just as he actually did, but had never met Bugsy in the interim and had never formed any further opinion about him, then it would be true that Elmer no longer believes that Bugsy is dangerous. If anyone can ever give up a formerly held belief about someone, then this would be a typical and central case of such an occurrence. But Elmer is actually in exactly the same state as this, save for the fact that he had met Bugsy in the interim and had formed an opinion about him at that time. Why should Elmer's former beliefs make any difference here? It is just his present doxastic state that we want to capture in specifying his disposition with respect to the information that Bugsy is dangerous. Elmer's present attitude toward this information involves something that ordinarily *constitutes* relinquishing a former opinion. Unless we find some appropriate way

to specify Elmer's withholding belief, we leave out of our account a very important element of Elmer's cognitive or doxastic state.

This seems to require us to say that Elmer does not believe that Bugsy is dangerous (or that Bugsy is not dangerous). But that contradicts something which we have also said, and which it appears we are required to say, concerning Elmer's befuddlement. Even during his soliloquy on June 1, Elmer steadfastly remained convinced of his friend's dangerousness. Thus, the facts of the matter in the story of Elmer's befuddlement seems to require us to say that Elmer still believes, at least since April 1, that Bugsy is dangerous, and they also seem to require us to say that Elmer no longer believes, as of June 1, that Bugsy is dangerous. Now, it sometimes happens that a story involves certain inconsistencies. For example, if the author of a series of mystery novels decides to alter some of the biographical facts concerning the detective who is the main character in all the novels (say, his birthdate), then stringing these novels together yields an inconsistent story. But nothing like this is the case with the story of Elmer's befuddlement. Clearly, the story is consistent. There is no logical reason why it could not be a true story. Perhaps some structurally similar befuddlement has actually occurred at some time in the history of intelligent life in the universe, or may yet occur at some time in the future.

Here, then, is the puzzle. Either Elmer believes that Bugsy is dangerous or he does not. Which is it? We seem to be required to say that Elmer does indeed believe that Bugsy is dangerous, for he remains convinced of his friend Bugsy's dangerousness. We also seem to be required to say Elmer does not believe that Bugsy is dangerous, for he now actively suspends judgment concerning the criminal's dangerousness. Yet we are logically prohibited from saying both together. How, then, are we to describe coherently Elmer's doxastic disposition with respect to the information that Bugsy Wabbit is dangerous? How can it be consistent for Elmer to believe that Bugsy is dangerous, on the one hand, and to withhold that belief, on the other?

The same puzzle can be stated with a different emphasis by focusing on the fact that on June 1 Elmer, in some obvious (but so far unclear) sense, changes his mind about whether Bugsy is dangerous. The change of mind is evident in Elmer's soliloquy. He suspends judgment where he used to have an opinion. Before that, at least since January, Elmer believed that Bugsy is dangerous. But there is also some obvious sense in which Elmer does not change his mind on June 1 concerning Bugsy, since he remains convinced of his friend's dangerousness. If we say, then, that Elmer continues to believe even on June 1 that Bugsy is dangerous, we fail to depict his change of mind. We represent him as believing on January 1 that Bugsy is dangerous, believing it still on

April 1, and believing it still even on June 1 after taking into account the further information from the FBI. There is nothing in all this about any change of mind. In order to express the fact that Elmer has changed his mind concerning Bugsy's dangerousness, we would like to say that Elmer believed on January 1 that Bugsy is dangerous but by the following summer believes it no longer (and also does not believe that Bugsy is not dangerous). However, we seem to be prevented from saying this; else we lie about Elmer's continued and unwavering conviction concerning his friend's dangerousness. How, then, do we express the important fact about Elmer that he has changed his mind concerning the question of Bugsy Wabbit's dangerousness and has withdrawn his former opinion?

7.2.3 Some Nonsolutions

It is worth mentioning some tempting nonsolutions to the puzzle. First, it is no solution to attribute an inconsistency solely to Elmer and his system of beliefs. At no point in the story does Elmer come to believe both that Bugsy is dangerous and that he is not dangerous. In fact, at no point does Elmer come to believe that Bugsy is not dangerous.[2] Moreover, even if an inconsistency is uncovered among Elmer's beliefs (e.g. that Bugsy $\neq$ Bugsy), that does not rescue us from the apparent meta-inconsistency concerning Elmer's beliefs to which we seem committed, for it is we who seem committed to saying both that Elmer believes and that he does not believe that Bugsy is dangerous.

One might try to avoid this inconsistency by looking to Elmer's idiolect. It may be argued that in Elmer's idiolect the syntactic sound and shape 'Bugsy Wabbit' represents not one but two different names, just as 'Aristotle' is ambiguous in the public language, functioning sometimes as a name referring to the celebrated philosopher of antiquity and sometimes as a name referring to the late shipping magnate. In Elmer's idiolect, one might argue, there are two distinct sentences corresponding to the syntactic string of the public language 'Bugsy is dangerous'. These two sentences might be represented formally by means of different subscripts on the name 'Bugsy'. We may then say, consistently, that Elmer still believes that Bugsy_2 is dangerous but no longer believes that Bugsy_1 is dangerous.

There are a number of difficulties with this attempt to solve the puzzle. In effect, it is an attempt to "reduce the problem to the previous case" (i.e., to cases like the 'Hesperus'-'Phophorus' and 'Cicero'-'Tully' examples, where there are distinct names for the same individual). I have been busy arguing that these cases are best treated in accordance with the modified naive theory. In any case, my main concern here is with the ability or inability, as the case may be, of the modified naive

theory to remove the puzzle. Accordingly, we should assume here that the modified naive theory is correct. On the modified naive theory, the information values of 'Bugsy$_1$' and 'Bugsy$_2$' are the same, viz., Bugsy Wabbit himself. Hence, if Elmer believes the information that Bugsy$_2$ is dangerous, ipso facto he also believes that Bugsy$_1$ is dangerous, for they are the same piece of information.

More important, this attempt to solve the puzzle does not even address the relevant question, for we are attempting to describe Elmer's doxastic state and we do not share Elmer's idiolect. In our idiolects, the name 'Bugsy Wabbit' is not ambiguous. We know what Elmer does not know: that there is only a single person of that name throughout the entire story. We use the expression 'Bugsy Wabbit' as a name for that individual. In our idiolects—and in the public language—the syntactic string 'Bugsy Wabbit is dangerous' unambiguously encodes a single piece of information, and the string 'Elmer believes that Bugsy is dangerous' contains the proposition that Elmer believes that information. The question is whether the latter string is true or false as it is used in our idiolects. We are not concerned with the truth value of other strings in other idiolects. We seem to be committed to saying that the relevant string, which is not ambiguous in our idiolects, is both true and false at the same time in our idiolects. But that cannot be right.

Again, one might try to avoid inconsistency by looking to other pieces of information. What Elmer believes, one might claim, is that his friend named 'Bugsy Wabbit' is dangerous, and what he suspends judgment about is whether the notorious criminal named 'Bugsy Wabbit' is dangerous. These are different pieces of information, and there is no contradiction in Elmer's believing the first but failing to believe the second.

This account is correct as far as it goes, but it also fails to address the problem. We are still left wondering what Elmer's doxastic disposition is with respect to the (singular) proposition that Bugsy is dangerous. Never mind what other propositions Elmer may believe or fail to believe; does he believe this proposition about Bugsy? There seems to be every bit as much reason to say that he does as there is to say that he does not, and vice versa. There are compelling reasons on both sides of the question, but taking either side seems utterly inadequate since it omits some critical element of Elmer's cognitive state. If we say that he continues to believe that Bugsy is dangerous, we must give some account of the sense in which Elmer changed his mind and now withholds belief. And, insofar as Elmer withholds belief, he withholds belief concerning this very same singular proposition, whatever other propositions he may also be withholding belief from (e.g., that the criminal named 'Bugsy' is dangerous).[3] Conversely, if we say that Elmer

does not believe that Bugsy is dangerous but suspends judgment, we need to give some account of the sense in which Elmer steadfastly retained his opinion formed shortly after meeting Bugsy, which is that he is dangerous. Surely nothing happened since then to deprive Elmer of this belief; today he would sincerely claim to retain this belief if asked.

Does the distinction between semantically encoded and pragmatically imparted information help solve the problem? It might be suggested that, strictly speaking, Elmer does believe the information that Bugsy Wabbit is dangerous, and that the temptation to say that Elmer does not believe this information results from a confusion of this information with further information that is only pragmatically imparted by utterances of 'Bugsy Wabbit is dangerous'. One difficulty with this attempt to solve the puzzle is that it is unclear exactly what information is supposed to be pragmatically imparted but not believed by Elmer. Apparently it is not the metalinguistic information that the relevant sentence is true, as in the previous examples, for whatever reason may be given for supposing that Elmer believes that Bugsy is dangerous may also be given for supposing that Elmer believes that the sentence 'Bugsy is dangerous' is true. Surely he *does* believe that this sentence is true, when it is understood as involving reference to his friend! A case would have to be made that utterances of the sentence pragmatically impart, say, the information that the notorious jewel thief named 'Bugsy Wabbit' is dangerous. Elmer suspends judgment concerning this information, and that would provide some account of why we are tempted to deny that Elmer believes that Bugsy is dangerous, even though he does have the belief. It is not clear, however, that such a case can plausibly be made. Surely not all utterances of the sentence, by Elmer or anyone else, pragmatically impart this information. Nor is it the case that typical utterances of the sentence, occurring in the context of the story, typically impart this information; some typical utterances, such as those made by Elmer (either aloud or to himself) on April 1, typically do not impart this information.

Even if the case can be made, the general account suggested here so far leaves it mysterious how it was decided that Elmer now believes that Bugsy is dangerous rather than that he suspends judgment. A full account of the situation must recognize somehow that, by summer, Elmer is in a doxastic state that would ordinarily *constitute* no longer believing that Bugsy is dangerous, but suspending judgment. As we have already seen, if Elmer had never met Bugsy in the interim, but had received the further information concerning Bugsy on June 1, just as he actually did, it would be false that Elmer continues to believe that Bugsy is dangerous. The doxastic state of mind that he is actually

in would constitute suspension of judgment; being in that state is all that would be required of him for it to be true that he has relinquished the relevant belief. That is just what giving up a belief is. Elmer's actual acquaintance with Bugsy, and Elmer's actual former opinions about him, should not stand in the way of this same state now constituting withheld belief and suspension of judgment.

The puzzle generated by the paradox of Elmer's Befuddlement is, to my mind, among the most difficult problems that arise in connection with propositional attitudes. Nevertheless, a certain extension of the modified naive theory contains the resources for solving this puzzle. The solution, sketched in the following chapter, suggests similar solutions to the new Frege Puzzle, and to many of the familiar philosophical problems that arise in this area.

8

Resolution of the Puzzles

8.1 Attitudes and Recognition Failure

To find a way out of the quandary generated by Elmer's befuddlement, I propose that we take seriously the idea that belief is a favorable attitude or disposition toward a piece of information or a proposition, and that we look more closely at the psychology and the logic of attitudes in general. Other favorable attitudes of the sort I have in mind are such states as that of liking ice cream, finding a certain piece of music pleasant, or loving someone. All of these are analogous in certain respects to belief of a proposition. For example, each seems to be a "standing" state—that is, it does not require the immediate presence of "occurrent" subjective experiences of approval or pleasure directed toward the object of the attitude at every moment while one is in the relevant state; it requires at most only occasional such occurrent experiences (typically when one is confronted with the object of the attitude). If de re belief is belief of a singular proposition, then de re belief about an external object is in certain important respects like loving someone. Both consist chiefly in the subject's adopting a certain favorable or positive attitude toward something external, or, in the case of de re belief, toward an abstract entity—a proposition—made up in part of something external. And in both cases, the adoption of the relevant favorable attitude can depend on the way in which the subject takes the object. If the subject does not recognize the object when it is encountered on different occasions, he or she may adopt the relevant attitude when the object is taken one way yet fail to adopt this attitude, and perhaps even adopt a corresponding unfavorable attitude (hatred or disbelief), when the object is taken another way.

Consider the following story, which is analogous to Elmer's Befuddlement. Suppose that Mrs. Jones does not realize it, but her husband leads a double life. By day he is the drab Mr. Jones, District Attorney, but by night he is Jones the Ripper-Offer (as he is called by the news media), a so-far-unidentified body snatcher who steals corpses from the city morgue. Mrs. Jones has faithfully loved her husband for many

years, but recently she has been intrigued and perversely fascinated by the macabre reports of Jones the Ripper-Offer. Stalking him out in the morgue, she eventually meets him but fails to recognize him as the very man she lives with and has lived with for many years. Unable to control her fascination, by April 1 she falls in love with Jones the Ripper-Offer. This bothers her deeply, since she has never fallen out of love with her husband, the DA. Emotionally, she is in that unfortunate state in which some people sometimes find themselves: being in love with (what she takes to be) two different individuals at the same time. Later that summer, her fascination with the demented body snatcher grows so overpowering that she retains no affection toward or emotional attachment to her husband whatsoever. She is now completely and entirely in love with the body snatcher, whom she still takes to be someone other than her drab husband.

Does Mrs. Jones now love Mr. Jones, alias Jones the Ripper-Offer, or doesn't she? Here again no simple 'yes' or 'no' answer by itself is satisfactory. Any attempt to describe Mrs. Jones's present emotional state with respect to Mr. Jones cannot rest only on the claim that she does love him (on the grounds that she remains in love with the grave robber), nor can it rest only on the claim that she does not love him (on the grounds that her emotional attitude toward him changed during the summer when she fell out of love with him). Somehow, both of these seemingly contradictory facts must be accommodated. But how?

Mr. Jones has two distinct personalities, two different guises. Under one of these guises, the happily married district attorney, Mrs. Jones once loved him but loves him no longer. Under his other guise, the demented grave robber, Mrs. Jones loved him before and loves him still. We do not normally speak of someone loving someone else under this or that guise. We say simply that *A* loves *B* or that *A* does not love *B*. The notion of *loving under a guise* is not the ordinary notion. But the case of Mrs. Jones's emotional attitude toward her husband is by no means a normal circumstance. In order to convey a complete picture of the situation, we must distinguish two ways in which Mrs. Jones can be in love with her husband. In April she loved him twice over, so to speak. By summer she loves him one way but not the other. We can decide to say that Mrs. Jones does love her husband by summer, in the absolute, nonrelativized sense of 'loves', since after all she does still love him in one of these two ways. But if we say only that Mrs. Jones loves Mr. Jones, we leave out of our description of Mrs. Jones's complex emotional attitude toward Mr. Jones the critical fact that she has fallen out of love with him and, in some obvious but unclear sense, loves him no longer. That is, if we allow ourselves only the ordinary, two-place, nonrelativized notion of *loves*, on which *A*

simply loves *B* or does not love *B*, the only thing that we can say about Mrs. Jones's present emotional attitude toward her husband—to wit, that she loves him—is seriously misleading at best, if not entirely and simply incorrect. It is only when we explicitly make the distinction between loving Jones qua her husband and loving him qua infamous body snatcher that we can coherently express the seemingly self-contradictory dual fact that Mrs. Jones has fallen out of love but also remains in love with a single man.

I do not claim that a three-place relativized notion of *loving qua*, or loving-in-a-certain-way, or loving-under-a-certain-guise, is philosophically clear or problem free. Surely it is not. What I do claim is that we have some grasp of this notion, and that it is clear in the present instance that Mrs. Jones loves Mr. Jones qua infamous grave robber (in this way, under this guise) but no longer loves him qua her husband (in that way, under that guise). The ordinary and familiar two-place notion of *A loving B* may then be identified with the existential notion of *A loving B in some way or other*, or *under some guise or other*, or *qua something or other*. At any rate, some such three-place notion of loving-in-a-certain-way or loving-under-a-certain-guise is required to capture all the relevant facts concerning Mrs. Jones's emotional state with respect to Mr. Jones, for in this special case the relevant three-place relation holds among the triple of Mrs. Jones, Mr. Jones, and one such third relatum by which Mrs. Jones is acquainted with Mr. Jones (whatever sort of thing this third relatum is, e.g. a guise), but fails to hold among Mrs. Jones, Mr. Jones, and another equally relevant third relatum. No account framed only in terms of a mere binary relation between Mrs. Jones and Mr. Jones can discriminate the relevant possibilities in this case and do justice to the relevant facts. Trying to get by with only the ordinary two-place notion of *loves* is like trying to specify whether an object is red by saying only whether it is colored, or like trying to convey whether 16 is odd or even allowing yourself only a predicate for being a composite (nonprime) integer.

8.2 Propositional Attitudes and Recognition Failure

Just as Mrs. Jones failed to recognize her husband on April 1 when she fell in love with him a second time, so Elmer failed to recognize Bugsy Wabbit on April 1 when Elmer formed the opinion for a second time that Bugsy is dangerous. But there is something else that Elmer failed to recognize: the information or proposition about Bugsy that he is dangerous.

The very idea of someone's failing to recognize a piece of information or a proposition can be somewhat mysterious. The phenomenon of

failing to recognize an individual person or material object is familiar. All of us have had the experience of running into someone who was familiar a long time ago but whose physical appearance has changed in the interim "beyond recognition," that is, to such an extent that we take him or her to be a perfect stranger. Many of us have had the converse experience of being taken for a total stranger by a past acquaintance. These are cases in which an individual goes unrecognized because of a significant objective change in physical appearance. In some cases, a change in physical appearance is induced intentionally for the precise purpose of preventing recognition. We call this 'taking on a disguise'. (Note the 'guise' in 'disguise'.) In our story of Elmer's Befuddlement, Bugsy disguised himself precisely so that Elmer would not recognize him. An object or an individual may also go unrecognized by a subject even though the object's appearance has not undergone any significant change. The subject may be situated with respect to the object in such a way as to prevent recognition, as when a familiar object is too far away for its distinguishing features to be discerned, or the subject's senses may be impaired. In such cases, although the object has not undergone any significant physical change in appearance, there is a change in what might be called its 'subjective appearance' with respect to the subject. This also occurs when the subject who is familiar with an object by having perceived some part of it encounters the same object by perceiving a different part of it, as may happen when a weary traveler passes the front of a building, inadvertently travels in a circle, and approaches the same building from the rear.

What all these cases of recognition failure have in common is that the object goes unrecognized by the subject because of a change in appearance—either objective appearance or subjective appearance. In both types of cases, the change of appearance may be a one-time-only affair, as with Bugsy, or the object may, so to speak, vacillate between two or more appearances, being regularly encountered in both or several of its appearances or guises, as in the case of Mr. Jones and the philosophical legend of the planet Venus. Now I am suggesting that Elmer has failed to recognize not only Bugsy himself but also a piece of information, or a proposition, concerning Bugsy. But propositions do not have *appearances*. We do not perceive propositions through the senses; propositions do not "appear" to us in the way that material objects do. If we "encounter" propositions at all, we do so by grasping or apprehending them—in an act of understanding, an act of the intellect, an act of thought or cognition. There is no notion of a proposition's *appearance*. Hence, there is no notion of a proposition's *changing its appearance*, and consequently there is no notion of *failing to recognize*

a proposition in the way that a subject can fail to recognize some individual because of a change in its objective or subjective appearance.

Of course, one may be said to "fail to recognize" a particular proposition—say, by failing to reidentify it as the very cornerstone of Mactaggart's philosophy of time, or as the very proposition to which the United States is said to be dedicated (the proposition that all men are created equal). In such cases, the proposition in question has been elected to some special sort of status, and it is this special status that the subject fails to impute to the proposition. In the same way, one might be said to "fail to recognize" a colleague in one's department by failing to think of him as, say, the world's foremost authority on the history of rock 'n' roll music from 1956 to 1959. This is not a case of mistaking him for a perfect stranger, nor is failing to think of the proposition that all men are created equal as the proposition to which the United States is said to be dedicated a case of mistaking that proposition for some other proposition in the way that I may mistake someone from my past to be a perfect stranger, someone unknown to me, someone with whom I am wholly unacquainted. Any proposition I can apprehend is a proposition that is fully known to me, in the relevant sense, for the only relevant sense in which one may be "acquainted with" a proposition is that one may fully apprehend it. If there is any sense to be made of a notion of "recognizing a proposition," analogous to recognizing a friend or acquaintance upon encountering him or her, it can only be this: that one fully apprehends the proposition. Once a proposition is fully apprehended, there is nothing relevant about the proposition that one is missing or failing to notice. To apprehend a proposition fully is to identify it in the fullest and most complete way that one can.

This objection to the notion of grasping but failing to recognize a proposition flows naturally from the traditional conception of the nature of propositions. In particular, it is the natural reaction of one who is thinking of propositions in accordance with the orthodox theory, that is, in accordance with the theories of Frege and Russell. But this is because one is in the grip of a faulty and misleading picture. On the Fregean conception, every piece of cognitive information, every "thought," is made entirely of things like concepts. (See the introduction.) To apprehend such a "thought" is, it seems, to be fully acquainted with it. There is no changing appearance, no superficial surface concealing the soul, no guise or veil of outward manifestation interceding between the subject and the thing-in-itself. To apprehend it is, as it were, to see through it, to see directly to its very soul. The same is true of a singular proposition whose only constituent other than things like concepts is a particular sensation or visual sense datum, an item

of "direct acquaintance." There is no "failing to recognize" a particular pain, for example by mistaking it for someone else's tickle. To have such a sensation or sense datum is to be acquainted with it in the fullest and most complete way possible. But the modified naive theory allows for propositions of a different sort: singular propositions involving external individuals and material objects as constituents. Clearly, the mode of apprehension for such propositions must be more complex than the mere grasping of pure concepts and the experiencing of wholly internal sensations. Apprehending such a proposition cannot be a wholly internal, mental act. The means by which one is acquainted with a singular proposition includes as a part the means by which one is familiar with the individual constituent(s) of the proposition. The mode of acquaintance by which one is familiar with a particular object is part of the mode of apprehension by which one grasps a singular proposition involving that object. For example, if one is familiar with some individual by having read his or her writings, then the reading of these writings is also part of the means by which one is acquainted with a singular proposition about that individual—say, that he or she had an unhappy childhood. One apprehends this proposition in part by having read the words written by the individual the proposition is about. If Elmer is familiar with Bugsy Wabbit by having interacted with him socially, then the social interaction with Bugsy is part of the means by which Elmer grasps the proposition about Bugsy that he is dangerous. Elmer apprehends this proposition in part by having interacted with part of it.

It is a large and difficult problem to specify exactly what sorts of modes of acquaintance with an object are sufficient to place one in a position to entertain singular propositions about that object. Must the mode of acquaintance be causal? Is any causal relation enough? (Consider the case of numbers and mathematical knowledge.) Is it enough simply to have heard the individual mentioned by name? Is it enough simply to be able to refer to the object? (Consider the shortest spy.) Is it enough simply to point at the object, without even looking to see what one is pointing at? Must one have some conception of what kind of thing the object is (a person, an abstract entity, etc.)? Can one have mistaken opinions about the object? How many? Does one have to know who the individual is, or which object the object is, in some more or less ordinary sense of 'know who' or 'know which'? Must one know some feature or characteristic of the object or individual that distinguishes it (or him or her) from all others? Is it sufficient simply to know some distinguishing feature or characteristic (i.e., is what Russell called 'knowledge by description' always enough)? It is not important

for the present purpose to have the answers to all of these questions, or even to any of them. What is important is to recognize that, whatever mode of acquaintance with an object is involved in a particular case of someone's entertaining a singular proposition about that object, that mode of acquaintance is part of the means by which one apprehends the singular proposition, for it is the means by which one is familiar with one the main ingredients of the proposition. This generates something analogous to an "appearance" or a "guise" for singular propositions. If an individual has a certain appearance, either objective or subjective, and through perceiving the individual one comes to have some thought directly about that individual—say, a thought that would be verbalized as 'Gee, is he tall'—then there is a sense in which the cognitive content of the thought may be said to have a certain appearance for the thinker since one of its major components does. This unorthodox conception of the nature of propositions and their apprehension thus allows for the possibility of a notion of "failing to recognize" a proposition by mistaking it for a new and different piece of information. If the subject happens to see the same tall man tomorrow without recognizing that it is the same man, and the subject happens to think 'Gee, is he tall', the subject's thought will have precisely the same cognitive content as the earlier thought, even though the subject does not recognize that this is so.

There is no reason why the modified naive theory should hold that the grasping of a piece of information places one in a position to "see through" the information, so to speak, and to recognize it infallibly as the same information encountered earlier in quite different surroundings under quite different circumstances. In fact, there is every reason to reject this idea.

8.3 Resolution

8.3.1 Elmer's Befuddlement

Now, whatever the necessary and sufficient conditions are for being in a position to entertain a singular proposition, it is clear that Elmer was in such a position on January 1, before he actually met Bugsy, when he first formed the opinion about Bugsy that he is dangerous. Elmer was an expert on Bugsy, well acquainted with his appearance and deeds through reports, photographs, tape recordings, and the rest; all these form a part of the means by which Elmer apprehends the proposition about Bugsy on January 1 that Bugsy is dangerous. Later, when Elmer meets up with Bugsy and forms for a second time the opinion that he is dangerous, Elmer apprehends this same proposition by entirely different means. His new mode of acquaintance with Bugsy,

and thereby with the proposition that he is dangerous, involves perceptions of a wholly new appearance. The proposition takes on a new guise for Elmer. In failing to recognize Bugsy, Elmer also fails to recognize the very proposition that he is dangerous. It is precisely for this reason that Elmer is able to form for the second time the opinion that Bugsy is dangerous without having ceased believing this very same piece of information. Elmer took his friend Bugsy to be someone other than the notorious jewel thief. Consequently, he took the information that he is dangerous, when it occurred to him on April 1, to be a different piece of information from the proposition about the jewel thief that he is dangerous (information that Elmer already believed). Elmer's problem stems from the fact that he takes the information that Bugsy Wabbit is dangerous to be two distinct and utterly independent pieces of information. He grasps it by means of two distinct appearances or guises; he takes it in two different ways. When he takes it in one way, Elmer does not recognize this piece of information as the same information that he also takes the other way. On June 1, Elmer adopts conflicting doxastic dispositions with respect to what he takes to be two different pieces of information but what is in fact a single proposition. On the one hand, Elmer has the appropriate favorable attitude toward this information; he is disposed to assent. On the other hand, he does not have an appropriate favorable attitude toward this information. It all depends on how Elmer takes the information.

How do we avoid this apparent contradiction? Does Elmer believe the relevant information, or doesn't he?

I have said that belief is a favorable attitude toward a piece of information, perhaps a disposition to inward assent or agreement. I have not said, however, that there must be a disposition to inward assent or agreement no matter how the information is taken. Elmer assents to the proposition that Bugsy is dangerous; he agrees with this information when he takes it as information concerning his friend. Hence, Elmer does believe this information. The fact that Elmer is no longer so disposed when he takes it as information concerning the notorious jewel thief does not entail that he has no disposition to assent to the proposition whatsoever. Indeed, he *has* such a disposition when he takes the proposition another way. This resolves the contradiction: Strictly speaking, Elmer does believe that Bugsy is dangerous, and it is strictly incorrect to say that he does not believe this, even after his change of mind on June 1.

We can still account for Elmer's change of mind with respect to the proposition that Bugsy is dangerous. When Elmer takes the information that Bugsy is dangerous as the information concerning his friend, he is continuously disposed to inward agreement since April 1. It is for

this reason that we say that Elmer continues to believe that Bugsy is dangerous. However, when Elmer takes the proposition to be one about the notorious jewel thief, he agrees with it on January 1 but by the following summer he is no longer so disposed. There is a certain way of taking the proposition that Bugsy is dangerous such that Elmer grasps the proposition by means of it but is no longer disposed to assent to the proposition when taking it that way. In this special sense, Elmer now *withholds belief*. Strictly speaking, this is not to say that he *fails* to believe. Nonetheless, Elmer manifests the central and most significant characteristic of giving up this belief so long as he takes the proposition to be one about the criminal, for then he is disposed to neither inward assent nor inward dissent, neither agreement nor disagreement, with respect to the relevant proposition. The only thing that prevents Elmer from failing to believe altogether is the fact that he happens to harbor a disposition to inward assent when he takes the proposition another way. This, at any rate, is how the modified naive theory can explain the sense in which Elmer may be said to "withhold belief". The fact one attempts to convey is just the fact that Elmer now lacks the appropriate favorable attitude or disposition when he takes the proposition in a certain contextually significant way.

I have argued so far as if belief may be analyzed in terms of a notion of *disposition to inward assent or agreement when taken in such-and-such a way*. It does not matter much whether this is the relevant notion, only that the modified naive theory is compelled to acknowledge some such ternary relation whose existential generalization coincides with the binary relation of belief. The matter can be put more formally as follows: Let us call the relevant ternary relation, whatever it is, *'BEL'*. It is a relation among believers, propositions, and something else (e.g. the relation of disposition to inward agreement when taken in a certain way), such that

(i) $\ulcorner A$ believes $p\urcorner$ may be analyzed as $(\exists x)[A$ grasps p by means of x & $BEL(A, p, x)]$,

(ii) A may stand in *BEL* to p and some x by means of which A grasps p, without standing in *BEL* to p and all x by means of which A grasps p,

and

(iii) $\ulcorner A$ withholds belief from $p\urcorner$, in the sense relevant to Elmer's befuddlement, may be analyzed as $(\exists x)[A$ grasps p by means of x & $\sim BEL(A, p, x)]$.[1]

In the special case of Elmer's Befuddlement, we initially seemed compelled to say both that Elmer believes that Bugsy is dangerous and

that Elmer does not believe that Bugsy is dangerous. The grounds for saying that Elmer does believe that Bugsy is dangerous are straightforward. Elmer formed this opinion on April 1 and has remained steadfastly convinced ever since. It is strictly incorrect, therefore, to say that Elmer does not believe that Bugsy is dangerous. How, then, do we express the other side of Elmer's doxastic state resulting from his recent change of mind? The specifics of the story do not allow us to say that Elmer believes that Bugsy is not dangerous, and so we are prevented from transferring the inconsistency from us to Elmer by saying that Elmer believes both that Bugsy is dangerous and that he is not dangerous. What, then, do we say to capture Elmer's apparent withheld belief, which we initially tried to capture by saying that he no longer believes that Bugsy is dangerous? The analysis in terms of *BEL* uncovers that there is yet a third position in which the negation sign may occur. What we are trying to say when we say, erroneously, that "Elmer does not believe that Bugsy is dangerous" is not

> $\sim(\exists x)$[Elmer grasps *that Bugsy is dangerous* by means of x & *BEL*(Elmer, *that Bugsy is dangerous*, x)]

(that is, it is not the case that Elmer believes that Bugsy is dangerous). This would saddle us with a contradiction. Nor is it

> $(\exists x)$[Elmer grasps *that Bugsy is not dangerous* by means of x & *BEL*(Elmer, *that Bugsy is not dangerous*, x)]

(that is, Elmer believes that Bugsy is not dangerous). This is straightforwardly false. Rather, it is

> $(\exists x)$[Elmer grasps *that Bugsy is dangerous* by means of x & $\sim$*BEL*(Elmer, *that Bugsy is dangerous*, x)]

(that is, Elmer withholds belief about Bugsy's being dangerous). This is at once true, compatible with Elmer's believing that Bugsy is dangerous, and *constitutive* of Elmer's change of mind. There is some relevant third relatum x such that on January 1 Elmer stands in *BEL* to the proposition that Bugsy is dangerous and x but by the following summer Elmer no longer stands in *BEL* to this proposition and x. As in the case of Mrs. Jones's complex emotional attitude with respect to her husband, alias Jones the Ripper-Offer, no attempt to describe Elmer's complex doxastic state with respect to the singular proposition about Bugsy Wabbit that he is dangerous can succeed using only the two-place notion of belief as a binary relation between believers and propositions. Without some relativized, ternary notion, and the resulting distinction between withholding belief and failure to believe, the attempt to describe Elmer's complex doxastic state with respect to the relevant

singular proposition breaks down. The only thing one can say using the binary notion of belief—to wit, that Elmer does believe the proposition that Bugsy is dangerous—is highly misleading at best. Thus, by casting singular propositions as objects of belief, the modified naive theory is compelled to acknowledge an analysis of belief as the existential generalization of some three-place relation *BEL* in order to uncover the appropriate position for the negation required by Elmer's change of mind in the face of his continued belief.

8.3.2 The New Frege Puzzle

This modified naive theoretical scheme for solving the problems posed by Elmer's Befuddlement points the way to a similar and related treatment of some of the other problems encountered earlier. Consider again the new and stronger version of Frege's Puzzle: An ancient astronomer-philosopher, who is an ideally competent speaker and thinker and a firm believer in the modified naive theory, unhesitatingly assents to 'Hesperus is Hesperus', but is not in the least disposed to assent to the sentence 'Hesperus is Phosphorus', even though he understands both sentences perfectly and, in fact, associates the very same proposition with each sentence. The explanation now available on the modified naive theory begins with the observation that the astronomer-philosopher does not recognize the proposition he attaches to the second sentence as the very same proposition he attaches to the first sentence, and firmly believes on logical grounds alone. When he reads and understands the sentence 'Hesperus is Phosphorus', he takes the proposition thereby encoded in a way different from the way in which he takes this same proposition when he reads and understands the sentence 'Hesperus is Hesperus'. He grasps the very same proposition in two different ways, by means of two different guises, and he takes this single proposition to be two different propositions. When he takes it as a singular proposition of self-identity between the first heavenly body sometimes visible in such-and-such location at dusk and itself, he unhesitatingly assents inwardly to it. When he takes it as a singular proposition identifying the first heavenly body sometimes visible in such-and-such location at dusk with the last heavenly body sometimes visible in so-and-so location at dawn, he has no inclination to assent inwardly to it, and may even inwardly dissent from it. His verbal assent and his refraining from verbal assent with respect to the two sentences are merely the outward manifestations of his inward dispositions relative to the ways he takes the proposition encoded by the two sentences. In the context of the new Frege Puzzle, this entails a rejection of the analogue of Frege's Law stated in terms of the verbal dispositions of

ideally competent speakers. Unlike Frege's Law, this analogue is not a truth of logic but an empirically false hypothesis.

The account of belief as the existential generalization of a ternary relation *BEL* was constructed around the modified naive theory's account of de re belief as a binary relation between believers and singular propositions (see the introduction), so that the modified naive theory could accommodate Elmer's complex cognitive state. The analysis makes room for the modified naive theory's claim that whoever believes that Hesperus is Hesperus also believes that Hesperus is Phosphorus, for whoever agrees inwardly with the singular proposition about the planet Venus that it is it, taking the proposition as an affirmation of self-identity about the first heavenly body sometimes visible at dusk in such-and-such location, stands in *BEL* to the proposition that Hesperus is Phosphorus and some x or other, and hence believes this singular proposition, even if he or she is not so disposed when this same proposition is taken some other way (e.g. as information concerning the last heavenly body sometimes visible at dawn in so-and-so location). It is part of the account that one who stands in the *BEL* relation to the information about Venus that it is it, together with some third relatum x by means of which he or she grasps this information, need not also stand in the *BEL* relation to this same information together with some further relatum y distinct from x and by means of which he or she also grasps the information.

8.4 Why We Speak the Way We Do

This aspect of the account yields another part (promised in section 6.2) of the explanation for the prevailing inclination to say—erroneously, according to the modified naive theory—that the ancient astronomer-philosopher does not believe that Hesperus is Phosphorus, and that Lois Lane is not aware that Superman is Clark Kent. The first part of the explanation was that most speakers, being insufficiently aware of the distinction between semantically encoded and pragmatically imparted information, will inevitably mistake information only pragmatically imparted by utterances of 'Hesperus is Phosphorus' (such as the information that the sentence is true) for part of the information content of the sentence, and hence will mistake the sentence 'The astronomer-philosopher believes that Hesperus is Phosphorus' for an assertion that the astronomer-philosopher believes this imparted information—information we know he does not believe. It was seen, however, that this explanation by itself cannot be the complete story, for, even when one takes care to distinguish semantically encoded and pragmatically imparted information, the astronomer-philosopher's fail-

ure to assent to the sentence 'Hesperus is Phosphorus', when he fully understands it and completely grasps the information thereby encoded, provides a compelling reason to suppose that he does not believe this information, and this reason is part of the original justification for denying that he believes that Hesperus is Phosphorus. The existential analysis of belief in terms of the ternary relation *BEL* reveals that this sort of evidence, compelling though it may be, is *defeasible*. When the astronomer-philosopher fails to assent verbally to 'Hesperus is Phosphorus', having fully understood the sentence, he also fails to assent mentally to the information thereby encoded, taking it in the way he does when it is presented to him through that particular sentence. He "withholds belief," in the sense defined earlier. This does not entail that he does not mentally agree with this information however he takes it. In the usual kind of case, one uniformly assents or fails to assent to a single piece of information, however it is taken, by whatever guise one is familiar with it. It is for this reason that failure to assent to a proposition when taking it one way—i.e., withholding belief—is very good evidence for failure to believe. But in this particular case it happens that the astronomer-philosopher is also familiar with the information that Hesperus is Phosphorus under its guise as a trivial truism, the way he takes it when it is presented to him through the sentence 'Hesperus is Hesperus'. Taking it this way, he unhesitatingly assents to it. Hence, he believes that Hesperus is Phosphorus, and his sincere denials constitute defeated, misleading evidence to the contrary. He "withholds belief," in the sense used here, but he also believes, in the sense used everywhere. To say that he does not is to say something false.

The true sentence 'The ancient astronomer believes that Hesperus is Phosphorus' may even typically involve the Gricean implicature, or suggestion, or presumption, that the ancient astronomer believes (his sentence for) the sentence 'Hesperus is Phosphorus' to be true and, under normal circumstances, would verbally assent to it if queried. Since he does not and would not, the implicatures of the sentence would also lead speakers to deny it, even though its literal truth conditions are fulfilled.

The reasons just given why we speak the way we do in cases of propositional recognition failure may still fail to get to the bottom of the problem. In attributing beliefs, we are stating whether the believer is favorably disposed to a certain piece of information or proposition. In the 'Hesperus'-'Phosphorus' and 'Superman'-'Clark Kent' cases, however, the believer in question is favorably disposed toward a certain singular proposition when taking it one way, but fails to recognize this proposition and is not favorably disposed toward it when it is en-

countered again. Since our purpose in attributing belief is to specify how the believer stands with respect to a proposition, we should, in these cases where the believer's disposition depends upon and varies with the way the proposition is taken, want to specify not only the proposition agreed to but also something about the way the believer takes the proposition when agreeing to it. The dyadic predicate 'believes' is semantically inadequate for this purpose; we need a triadic predicate for the full *BEL* relation, which the belief relation existentially generalizes. But there may be no such predicate available in the language. Even if such a predicate is available, it may be inordinately long, or cumbersome, or inconvenient. We are accustomed to speaking with the dyadic predicate, 'believes', and we mean to continue doing so even in these problem cases. How, then, do we convey the third relatum of the *BEL* relation?

In the case of Elmer's believing that Bugsy Wabbit is dangerous, the sentence used to specify the information content of Elmer's belief, 'Bugsy Wabbit is dangerous', is itself understood by Elmer, though Elmer understands the sentence in two different ways. He mistakes the sentence to be semantically ambiguous. As one might say, he takes the single sentence to be two different sentences. This is unlike the 'Hesperus'-'Phosphorus' and 'Superman'-'Clark Kent' cases. In these cases, the two ways in which the believer takes the relevant proposition are associated, respectively, with two different sentences, either of which may be used in specifying the content of the belief in question. The ancient astronomer agrees to the proposition about the planet Venus that it is it when he takes it in the way it is presented to him through the logically valid sentence 'Hesperus is Hesperus', but he does not agree to this same proposition when he takes it in the way it is presented to him through the logically contingent sentence 'Hesperus is Phosphorus'. The fact that he agrees to it at all is, strictly speaking, sufficient for the truth of both the sentence 'The astronomer believes that Hesperus is Hesperus' and the sentence 'The astronomer believes that Hesperus is Phosphorus'. Though the sentences are materially equivalent, and even modally equivalent (true with respect to exactly the same possible worlds), there is a sense in which the first is better than the second, given our normal purpose in attributing belief. Both sentences state the same fact (that the astronomer agrees to the singular proposition in question), but the first sentence also manages to convey *how* the astronomer agrees to the proposition. Indeed, the second sentence, though true, is in some sense inappropriate; it is positively misleading in the way it (correctly) specifies the content of the astronomer's belief. It specifies the content by means of a 'that'-clause that presents the proposition in the "the wrong way," a way of taking the proposition with

respect to which the astronomer does not assent to it. This does not affect the truth value of the second sentence, for it is no part of the semantic content of the sentence to specify the way the astronomer takes the proposition when he agrees to it. The 'that'-clause is there only to specify the proposition believed. It happens in the 'Hesperus'-'Phosphorus' type of case that the clause used to specify the believed proposition also carries with it a particular way in which the believer takes the proposition, a particular x by means of which he or she is familiar with the proposition. In these cases, the guise or appearance by means of which the believer would be familiar with a proposition at a particular time t were it presented to him or her through a particular sentence is a function of the believer and the sentence. Let us call this function f_t. For example, $f_t(x,S)$ might be the way x would take the information content of S, at t, were it presented to him or her through the very sentence S. In the case of the ancient astronomer, we have

(7) *BEL*[the astronomer, that Hesperus is Hesperus, f_t(the astronomer, 'Hesperus is Hesperus')]

and

(8) *BEL*[the astronomer, that Hesperus is Phosphorus, f_t(the astronomer, 'Hesperus is Hesperus')],

but not

(9) *BEL*[the astronomer, that Hesperus is Hesperus, f_t(the astronomer, 'Hesperus is Phosphorus')]

and not

(10) *BEL*[the astronomer, that Hesperus is Phosphorus, f_t(the astronomer, 'Hesperus is Phosphorus')].

The quasi-symbolizations 7 and 10 reveal that, though one cannot be explicit about the particular third relatum involved in the *BEL* relation using only the dyadic predicate 'believes', one can, so to speak, "fake it" by using as a 'that'-clause a sentence that determines the third relatum in question. If one existentially generalizes on the third argument place in all of sentences 7–10, the first and the fourth, unlike the second and the third, typographically retain all that is obliterated by the variable of generalization—all, that is, but the functor 'f_t' and the quotation marks around and recurrence of its second argument. One can exploit this feature of the sentence 'The astronomer believes that Hesperus is Hesperus' to convey the third relatum of *BEL*. The 'that'-clause, whose semantic function is simply to specify the content of the astronomer's belief, is also used here to perform a pragmatic function involving an

autonomous mention-use of the clause. This is the closest one can come to saying by means of the dyadic predicate what can, strictly speaking, be said only by means of the triadic predicate. To borrow Wittgenstein's terminology, one *shows* using 'believes' what one cannot *say* by its means alone.

Since it is our purpose in this case to convey not only what the astronomer agrees to but also how he takes what he agrees to when agreeing to it, the belief attribution 'The astronomer believes that Hesperus is Phosphorus' may typically involve the false (further) implicature (or suggestion, or presumption) that the astronomer agrees to the proposition that Hesperus is Phosphorus when he takes it in the way it is presented to him through the very sentence 'Hesperus is Phosphorus'. If we allowed ourselves the full triadic predicate, we could cancel the implicature without explicitly specifying the third relatum by uttering something like the following:

> The astronomer believes that Hesperus is Phosphorus, although he does not agree that Hesperus is Phosphorus when he takes this information the way he does when it is presented to him through the very sentence 'Hesperus is Phosphorus'.

The second conjunct here—the cancellation clause—is meant to take the sting out of the first conjunct, and the conjunction taken as a whole remains perfectly consistent. However, since the sentence that determines (via the function f_i) the way the astronomer takes the information when agreeing to it is readily available, it is easier and equally efficacious simply to retain the dyadic predicate 'believes' and to deny the literally true but misleading belief attribution 'The astronomer believes that Hesperus is Phosphorus' while asserting an equally true but not misleading attribution. Denying the misleading attribution is the closest one can come, using only the dyadic predicate, to denying proposition 9 (= proposition 10). Hence we are naturally led to say things like 'The astronomer believes that Hesperus is Hesperus, but he does not believe that Hesperus is Phosphorus'. We speak falsely, but the point is taken, and that is what matters. So it is that the modified naive theory, properly extended to acknowledge that believers may fail to recognize the singular propositions they embrace, predicts the sort of usage in propositional attitude discourse that we actually find where propositional recognition failure is involved.[2]

9

The Orthodox Theory versus the Modified Naive Theory

9.1 Semantics and Elmer's Befuddlement

The orthodox theory has the consequence that it is true that one can believe that Hesperus is Hesperus without believing that Hesperus is Phosphorus. This is a merit of the theory because it conforms with ordinary usage. The hard datum here is that this is the way we speak. But this datum by itself has no eristic value, for, once it is acknowledged that one's disposition with respect to a proposition may depend on how the proposition is taken, this hard datum—the fact that we say things like 'Lois Lane does not realize that Clark Kent is Superman' and 'One can believe that Hesperus is Hesperus without believing that Hesperus is Phosphorus'—is accommodated by the modified naive theory as well as by the orthodox theory. Since both theories accommodate the datum, we cannot use the datum to decide between the two theories. To turn this datum into the oft-heard objection against the modified naive theory that such pronouncements are indeed true is to beg the question against the modified naive theory. Pre-theoretically, all we have is that we speak this way. The truth values of our pronouncements are not pre-theoretic data. The orthodox theory appears to assert that what we say when we speak this way is true; the modified naive theory asserts that it is false. Only after we have decided on one theory over the other can we determine the truth value of our pronouncements. It is the theory that tells us whether what we say is true, not the other way around.

There is still some remnant of Frege's Puzzle remaining on the sort of account I have advocated here. I have argued that belief must be the existential generalization of a ternary relation *BEL* among believers, propositions, and something else—something which varies with the different ways in which a believer may be familiar with a proposition. But I have given at most only a vague sketch of what this relation *BEL* may be, suggesting as one candidate the relation of assenting inwardly, or being disposed to assent inwardly, to a proposition when taking it in a certain way. The account remains incomplete until more is said

about this. In the meantime, some challenging "questions which are not altogether easy to answer" (see section 1.1) arise: What exactly is the ternary relation *BEL*, and what is the nature of the sort of thing that serves as its third relatum? If *A*'s believing *p* consists in there being something *x* such that *A*, *p*, and *x* stand in the *BEL* relation, what is this extra something *x*? Is it a way of taking the proposition? Is it a mode of presentation of the proposition? Is it perhaps another proposition, or a sentence in the language of thought? Is it a "mental file"? What sort of thing is it, and how are such things individuated?

In fact, it is evident that the things that serve as third relatum for the *BEL* relation must be similar in some respects to Fregean senses.[1] Does this mean that the augmented version of the modified naive theory advocated here is simply a version of the orthodox, Fregean theory in a different dress? Definitely not. On the Fregean theory, senses are integral to the semantic nature of sentences. They make up the pieces of cognitive information, or "thoughts," that sentences encode. On the theory advocated here, the objects that serve as third relatum for the *BEL* relation, whatever sort of things they are, are entirely separable from the semantic nature of the relevant sentence encoding the second relatum of *BEL* (though it may turn out that the objects that serve as third relatum for *BEL* are characterizable using semantic notions). Sentences are devices for encoding information. The pieces of information they encode are propositions, often singular propositions. The semantics for a sentence like 'Socrates is wise' need treat only of the singular proposition about Socrates that he is wise; there is no need to consider the *x* by means of which some thinker may be familiar with that piece of information as long as one's concern is chiefly semantical (having to do with the cognitive information content of the sentence) and not psychological.[2]

The difference between this version of the modified naive theory and the Fregean theory, and also the theory of Russell, is brought out dramatically in their differing treatments of Elmer's Befuddlement. I have argued that, on the modified naive theory, strictly speaking, Elmer believes that Bugsy Wabbit is dangerous even after receiving the further information from the FBI on June 1. The orthodox theory delivers just the opposite verdict—at any rate, it *can* deliver the opposite verdict, and on the most plausible construal it does. Consider, for instance, how a Fregean would answer the question of whether Elmer believes that Bugsy Wabbit is dangerous. On the Fregean theory, the name 'Bugsy Wabbit' as used by us has a certain sense, and this sense is partly constitutive of the belief we ascribe to Elmer in uttering the sentence 'Elmer believes that Bugsy Wabbit is dangerous'. Of course, on the Fregean theory the name also has a sense—in fact, two distinct

senses—as used by Elmer. For Elmer, according to the Fregean, the name is strictly ambiguous, standing at once for the possibly dangerous jewel thief named 'Bugsy Wabbit' and also for his definitely dangerous friend named 'Bugsy Wabbit'. As we saw earlier, however, for us the name is univocal. This must be so even on the Fregean theory—at any rate, the name is most plausibly construed on the Fregean theory as univocal as used by us, for we know what Elmer does not know: that the jewel thief named 'Bugsy Wabbit' is the same person whom Elmer has befriended, having failed to recognize him as the jewel thief. On the Fregean theory, the name 'Bugsy Wabbit' expresses for us a sense like that of 'the crafty jewel thief who has tricked Elmer into believing him to be someone other than the very criminal that Elmer is pursuing', or 'the notorious criminal whom Elmer is pursuing but has befriended, having failed to recognize him for who he is', or something along these lines. Now, on the Fregean theory, when we report Elmer's beliefs using the name 'Bugsy Wabbit', we use the name to refer to *our* sense—at any rate, that is the most plausible way of construing us (unless we explicitly or implicitly signal that we are using the name with some sense that is nonstandard for us). On the Fregean theory, then, if we utter the sentence 'Elmer believes that Bugsy Wabbit is dangerous' we attribute to Elmer something like the belief that the crafty jewel thief named 'Bugsy Wabbit' whom Elmer has befriended, having failed to recognize him, is dangerous. Of course, Elmer has no such belief. Hence, the Fregean theory must claim that Elmer does not believe that Bugsy Wabbit is dangerous. At any rate, the Fregean theory must claim that, in at least one plausible but literal sense, Elmer does not believe that Bugsy Wabbit is dangerous. A similar situation obtains on the theory of Russell. The modified naive theory, as partially developed here, claims that, in the only literal sense, Elmer believes that Bugsy Wabbit is dangerous (though, in one sense of 'withholds belief', he also withholds belief concerning whether Bugsy is dangerous). Thus, the modified naive theory and the orthodox theory are diametrically opposed on one of the central issues with which both theories are concerned. This demonstrates that the augmented version of the modified naive theory urged here is no mere variant of the orthodox theory.

9.2 *Quantifying In*

The situation may appear to end in a stalemate, with both of two incompatible theories accommodating the data and no way to choose between the two. But that is not so. In the first place, the modified naive theory is the natural and compelling result of a preliminary philosophical investigation into the nature and structure of information con-

tent. Even Frege and Russell, who argued in opposition to the naive theory, came to the philosophy of language with an initial predisposition toward something like the naive theory. The modified naive theory has a prima facie claim on our endorsement; it must be refuted before it is abandoned. If all other things are equal, the modified naive theory is to be preferred over its rivals. The orthodox theory was invented because it was believed that the modified naive theory falters over attitude contexts. If I am correct that this belief is erroneous, we no longer have that reason to turn away from the original view.

An analogy from the philosophy of perception is helpful here. The natural preliminary theory of perception is that perceiving is a relation between a subject and an external object. Seeing an apple is an experiential "getting in touch" with an external reality. Let us call this the *naive theory of perception*. It is more or less the common-sense view. Now suppose that some clever philosophers were to give the following argument against the naive theory of perception: Clark Kent is the same individual as Superman. But when Lois Lane looks across her desk at the mild-mannered reporter, she sees Clark Kent and does not see Superman. It is only when Lois looks at the red-caped man of steel in blue tights with the letter 'S' on his chest that we should say that she sees Superman, and then we should say that she does not see Clark Kent—even though Clark Kent is Superman. According to these philosophers, perceiving is not a relation between a subject and an external object (otherwise Lois sees Superman if and only if she sees Clark Kent), but a relation between a subject and an internal object. Seeing an apple is having a visual apple-image. In some cases the image may not even represent an external reality. If one is hallucinating, for example, one may see an apple when there is no apple there. Perceiving per se is not a "getting in touch" with an external reality. Fortunately, the objects of perception—images, sensations, "ideas"—typically fit with the external surroundings, but that is beside the point. Perception per se is a wholly internal matter.

This theory of perception carries over into a doctrine in logic and semantics. According to these philosophers, perceptual contexts like 'Lois Lane sees ____' are oblique; 'Lois Lane sees Superman' asserts a relation between Lois Lane and her visual Superman-image. These philosophers acknowledge (reluctantly) that there is a notion of what might be called 'relational perception', which arises from quantification into perceptual contexts, e.g. 'Lois Lane sees something'. Once such locutions are acknowledged, it must also be allowed that there is indeed a notion of perception as a relation to this or that particular external thing. This notion of de re perception arises from sentences like 'Lois sees someone who is Superman'. But, say these philosophers, it is

obscure and puzzling just what de re perception amounts to, if anything. If one wishes to speak of de re perception as a relation between a subject and an external object, this de re relation must be *defined* as the relative product of the real perceiving relation between a subject and an internal "idea," and some representation relation (perhaps the relation of *fit*) between the "idea" and the external object. The ordinary notion of perception, which we typically invoke when we say things like 'Lois Lane sees Superman' and 'John sees the apple', they contend, is not this perplexing notion of de re perception.

Let us call this the *sophisticated theory of perception.* It is not to be confused with indirect realism. Indirect realism is the doctrine that one perceives external objects only indirectly, by directly perceiving internal objects that represent the external ones. Indirect realism is a form (proper extension) of the naive theory of perception. The naive theory merely asserts that external objects are in the range of the perceiving relation. It is perfectly compatible with this that external objects are in the range of the perceiving relation *by virtue of* internal objects which are already in the range of the perceiving relation and which represent the external objects. Since indirect realism is a form of the naive theory of perception, it is in fact incompatible with the sophisticated theory. The latter denies that external objects can enter into the range of the perceiving relation. (Recall that, on the sophisticated theory, the so-called de re perceiving relation is not the ordinary perceiving relation.)

The dispute presented here between the naive and the sophisticated theory of perception is parallel in a number of ways to the dispute between the modified naive and the orthodox theory of information content. Some morals may be drawn by pursuing the various analogies. For present purposes, the important point to recognize is that the burden of proof falls squarely on those who reject the naive theory of perception. The naive theory of perception is the natural, common-sense view. The sophisticated theory of perception is sophisticated. It was invented as a reaction to a certain philosophical argument that purports to refute the common-sense view. The sophisticated theory is built around the anomalous fact that we sometimes say things like 'Lois Lane saw Clark Kent; she did not see Superman'. This fact is then adduced as evidence in favor of the sophisticated theory. If it is discovered that the naive theory of perception, or a non-ad-hoc extension of it, also predicts this same data, the naive theory is thereby restored to its rightful place as the preferred theory. So much the better for the original theory if it can be extended in such a way as to incorporate some of the features of the rival, sophisticated view that give the latter its philosophical appeal.

More important, the truth of the matter is that the orthodox theory of information content does not accommodate all the data as well as the modified naive theory. The fact that we speak a certain way in propositional-attitude contexts where ignorance of an identity is involved is one datum that the orthodox theory appears to accommodate (but see below). There are still further data on which to test a theory of information content, and here the orthodox theory faces a number of serious difficulties that do not arise on the modified naive theory. I have already cited some of these difficulties in arguing in section 5.1 against the orthodox theory's identification of information value with conceptual content: twin-earth arguments, the argument from subjectivity of conceptual content, the argument from error in conceptual content, the modal arguments, and the epistemological arguments. Insofar as one views the situation as a contest between two main competitors, the modified naive theory of information content and the orthodox theory, these arguments lend further support to the modified naive theory.

In fact, the parallel with the so-called sophisticated theory of perception brings out yet a further, related difficulty with the orthodox theory, concerning the intelligibility of quantification into modal and propositional attitude contexts. De re belief locutions of the form ⌜there is some ψ which a believes to be $\overline{\phi}$⌝, or ⌜$\overline{a}$ believes of b that it is $\overline{\phi}$⌝, are used very widely in psychology (for example, in explaining behavior), and may well be indispensable to that discipline, as well as others. De re propositional-attitude locutions are remarkably pervasive in ordinary, everyday propositional-attitude discourse. Modal analogues, such as ⌜$\overline{b}$ is something which has to be $\overline{\phi}$⌝, are also not uncommon in ordinary discourse. (The modality contained in phrases like 'has to' need not be the philosopher's "metaphysical modality" or "modality in the strictest sense".) On the modified naive theory, quantification into modal and propositional-attitude contexts is logically and semantically straightforward: ⌜it is necessary that x is $\overline{\phi}$⌝, under an assignment of an individual i as value for the variable 'x', attributes (the relevant sort of) necessity to the singular proposition about i that it is ϕ. Similarly, ⌜$\overline{a}$ believes that x is $\overline{\phi}$⌝, under an assignment of i to 'x', attributes belief of the singular proposition about i that it is ϕ. (See the introduction.) On the orthodox theory, however, it is quite mysterious just what these de re locutions amount to. On that theory, quantification into modal or propositional-attitude contexts is something like quantification "into" quotation marks. A nonstandard interpretation is called for. The most natural way of interpreting quantification into modal or propositional-attitude discourse within the orthodox theory is to interpret ⌜it is necessary that x is $\overline{\phi}$⌝ as ⌜$(\exists\alpha)$(α determines x and it is

necessary that α is $\bar{\phi}$ and to interpret $\ulcorner a$ believes that x is $\bar{\phi}\urcorner$ as $\ulcorner(\exists\alpha)(\alpha$ determines x and a believes that α is $\bar{\phi}\urcorner$, where 'α' ranges over Fregean individual concepts (singular term senses), and the word 'that' is regarded as a proposition-term-forming sentential operator that functions in a manner analogous to quasi-quotation marks. But these interpretations cannot be correct, since they would render the de re locutions trivial and rarely if ever false. How, then, shall we understand these locutions? This is the alleged problem of quantifying in. There have been several attempts to provide a proper analysis or definition for these de re locutions in terms acceptable to the orthodox theory. Unfortunately, none of the proposed definitions is uncontroversial, clear, and clearly able to handle a wide range of cases. But why should there be any problem here in the first place? Quantification into modal and propositional-attitude contexts should be no more enigmatic than quantification into perceptual contexts. Either appears mysterious only when one is in the grip of a certain sort of theory. The fact that quantification into modal and propositional-attitude contexts is logically and semantically problematic on the orthodox theory points to a flaw in the theory and reveals another aspect in which the modified naive theory is superior to the orthodox theory. The "problem of quantifying in" is, at bottom, a pseudo-problem that arises from thinking in accordance with a false theory.[3]

9.3 *Propositional-Attitude Attributions*

The orthodox theory also faces serious difficulties in propositional-attitude attributions where closed proper names, demonstratives, or single-word indexicals are involved. On the usual formulations of the orthodox theory, the locution $\ulcorner a$ believes that b is $\bar{\phi}\urcorner$ attributes to the referent of a belief of a general proposition or "thought," made up in part of the conceptual content, or sense, of the singular term b. But, as was noted in subsection 5.1.3, the conceptual content attached to a proper name, as used with a particular reference, varies significantly from speaker to speaker. The same is true of demonstratives and single-word indexicals such as 'I', 'you', and 'here'. I do not know what conceptual content Plato attached to the ancient Greek version of the name 'Socrates'. In fact, about the only thing I do know concerning Plato's concept of Socrates is that it surely does not coincide exactly with mine. In fact, it is extremely unlikely that Plato should even have had my Socrates-concept in his repertoire of concepts, or that I should have his in mine. I cannot use the name 'Socrates' in Plato's sense, attaching to it Plato's conceptual content for the ancient Greek version of 'Socrates'. I use the name with my own conceptual content. If the

singular term b occurring in $\ulcorner a$ believes that b is $\phi\urcorner$ is a proper name, according to the usual formulations of the orthodox theory it is used there to refer to the speaker's sense for the name. It is even more obvious in the case where b is the indexical 'I' that it must be used in the sense given to it by the speaker rather than that given to it by the referent of a. The conceptual content which the subject of the attribution (the referent of a) happens to attach to b is entirely irrelevant to the attribution. Hence, according to the usual version of the orthodox theory, if I utter the sentence 'Plato believed that Socrates is wise' I attribute to Plato a belief made up in part of my concept of Socrates, the sense I attach to the name 'Socrates'. Almost certainly, Plato had no such belief. (He could not have believed, for example, that the ancient Greek philosopher I first learned about at Lincoln Elementary School in Torrance, California, is wise; it is part of the human predicament that one's concept of the Other tends to be constructed around one's concept of the Self.) The orthodox theory yields the result that, in principle, Lois Lane could believe that Clark Kent is mild-mannered without believing that Superman is, but the theory does so only at the cost of misrepresenting what the belief that Clark Kent is mild-mannered *is*—misrepresenting it to such an extent that, according to the theory, Lois does not have it.[4] According to the story, however, Lois does believe that Clark Kent is mild-mannered. Even with respect to their respective accounts of propositional-attitude discourse, then, the modified naive theory, properly extended to take account of the fact that one may fail to recognize a proposition that one embraces, is superior to the orthodox theory.

9.4 Concluding Remarks

The major problem remaining for the sort of theory I have advocated here is to provide a more complete account of the things corresponding to propositional recognition failure, the things that serve as third relatum for the *BEL* relation. This is by no means a trivial problem. The fact that some challenging questions are left unanswered does not mean that no progress has been made, however, for the remaining questions are not the same as those posed by Frege's original puzzle, and it seems likely that the newer questions are answerable. Propositional recognition failure is only a special case of the general phenomenon of recognition failure. Failure to recognize a proposition is often the result of failure to recognize some component of the proposition. Typically, the unrecognized proposition is singular and the unrecognized component is an individual that the proposition is directly about, but in some cases the unrecognized component may be something other than an indi-

vidual, e.g. a natural kind. A closer examination of the general phenomenon of recognition failure should deepen our understanding of the problems raised by Frege's Puzzle and the apparent failures of substitutivity in propositional-attitude contexts. To use a distinction of Kripke's, the account I have offered may not be a complete *theory* of information content and propositional attitudes, but I believe it yields a better "picture" of what is going on in connection with Frege's Puzzle and related problems than that given by the received view. My hope is that seeing the general problems posed by phenomena like Frege's Puzzle in the light of what has been said here will help to diffuse the idea that these problems pose a threat to the modified naive theory and to reshape the problem into something amenable to a final solution consistent with the modified naive theory. To my mind, however, an important aspect of Frege's Puzzle remains unsolved.

In addition to Frege's Puzzle and the related difficulties involving propositional-attitude discourse, the other major sources of objections to the modified naive theory have traditionally been the apparent existence of true negative existentials involving nonreferring names and the more general problem of the truth value and information content of sentences involving nonreferring names. Though my concern in the present book is exclusively with the former source of objections, a complete defense of the modified naive theory would require a separate discussion of the latter source. My defense of the modified naive theory against the former objections has been constructed essentially from two central ideas: the distinction between semantically encoded and pragmatically imparted information, and the explicit acknowledgment of something like ways in which a proposition is taken or guises by which one may be familiar with a proposition. It is important to recognize that either or both of these ideas might be effective also in removing the objections arising from nonreferring names. We have already seen that something can pragmatically impart information even if there is no piece of information that it semantically encodes. Depending on what sort of thing serves as the third relatum of the *BEL* relation, it might also turn out that there are things of that sort (e.g. ways of taking a proposition) to which there does not correspond any piece of information (e.g., such that there is no proposition which it is a way of taking). Also, a pair of sentences that differ only in containing different nonreferring names, demonstratives, or other single-word singular terms, and are otherwise exactly the same, may semantically encode the very same information, though each presents its information content to a particular speaker by means of a different way of taking it. Any two such sentences pragmatically impart different information. Some or all of these facts concerning sentences involving nonreferring

names or other single-word singular terms may be directly relevant to the philosophical problems that arise on the modified naive theory in connection with such sentences. A final verdict on the modified naive theory's ability to handle such sentences must await further investigation of the general problem. It is premature to dismiss the modified naive theory on the basis of these apparent difficulties.

Appendix A
Kripke's Puzzle

It is a simple matter to extend the account given of Elmer's Befuddlement to familiar problem cases, including Quine's case in which Ralph believes that Ortcutt is a spy, Mates's embedded propositional attitude contexts, Castañeda's examples concerning belief about oneself, and Kripke's puzzle in "A Puzzle About Belief." The modified naive theory's solution to Kripke's puzzle is briefly presented here for illustration.

Kripke's puzzle can be obtained from Elmer's Befuddlement by a slight modification of the story. Let everything happen the way it happens in the story of Elmer's Befuddlement up to, but not including, April 1. Thus, Elmer has already decided that Bugsy Wabbit, the notorious jewel thief, is dangerous, and he has met up with Bugsy, but he believes that this Bugsy is not the man he seeks. Now, suppose that Bugsy's transformation is so complete that Elmer is deceived even about Bugsy's dangerousness. On April 1, Elmer says to himself: "It's funny that my new friend should be named 'Bugsy Wabbit', since he is nothing at all like the criminal of that name. In particular, the criminal Bugsy Wabbit is a dangerous fellow, but my friend Bugsy Wabbit is perfectly harmless, not in the least bit dangerous." Does Elmer believe that Bugsy is dangerous, or does he believe the opposite? As with the previous examples, no answer seems satisfactory; hence the puzzle. (See note 1 to chapter 7.)

In constructing his puzzle, Kripke relies on instances of some (perhaps weakened) version of the following principle schema, which he calls 'the disquotation principle': *If a normal English speaker, on reflection, sincerely assents to 'S', then he believes that S*, where the substituends for '*S*' are "appropriate standard English sentences lacking indexical or pronominal devices or ambiguities."[1] Some commentators have urged solutions to Kripke's puzzle that involve rejecting Kripke's disquotation principle. It has been emphasized in the present book, however, that at least some version of this disquotation principle is unobjectionable; it is no solution to Kripke's puzzle to reject this principle. Kripke remarks that, "taken in its obvious intent, after all, the principle appears to be

a self-evident truth."[2] What makes the principle self-evident is that it is a corollary of the traditional conception of belief as inward assent to a proposition. Sincere, reflective, outward assent (qua speech act) to a fully understood sentence is an overt manifestation of sincere, reflective, inward assent (qua cognitive disposition or attitude) to a fully grasped proposition. In fact, an alternative version of Kripke's puzzle can be generated by means of a different and obviously unobjectionable disquotation principle concerning assertion in lieu of belief. In our example, Elmer thinks out loud. Thinking about the criminal, Elmer announces 'Bugsy is dangerous'; thinking about his friend, he announces 'Bugsy is not dangerous'. An appropriately analogous disquotation principle conditionally linking assertion to utterance of a sentence is very difficult to deny.[3] It seems, then, that Elmer has said (asserted) both that Bugsy is dangerous and that he is not. Does Elmer contradict himself, then? It seems incorrect to say that he does, since Elmer understands what he utters and knows what he is asserting when he utters the first sentence, and he also knows what he is asserting when he utters the second sentence. Yet even if he is a master logician he will not see any contradiction in the joint assertion. How can that be? An answer to this question must preserve the disquotation principle linking utterance and assertion. The modified naive theory, as developed here, provides at least a sketch of such an answer.

Kripke does not himself take an official position with respect to the question of the correct solution to his puzzle, but he seems to suggest that it may be somehow indeterminate or neither true nor false in this sort of case to describe Elmer as believing (or asserting) that Bugsy is dangerous, and that it may be indeterminate or neither true nor false to describe Elmer as believing that Bugsy is not dangerous: "The situation of the puzzle seems to lead to a breakdown of our normal practices of attributing belief. . . . [This is] an area where our normal practices of interpretation and attribution of belief are subjected to the greatest possible strain, perhaps to the point of breakdown. So is the notion of the *content* of someone's assertion, the *proposition* it expresses."[4]

From the point of view of the theory defended here, no such conclusion with respect to Kripke's puzzle is warranted. The simple fact is that Elmer believes (asserts) both that Bugsy is dangerous and that Bugsy is not dangerous, and it is simply false to say that he fails to believe (assert) either. With respect to both propositions—that Bugsy is dangerous and its negation—Elmer inwardly assents to (asserts) the proposition taking it a certain way. Before Elmer met Bugsy he formed the opinion that Bugsy Wabbit is dangerous; he believed the singular proposition about Bugsy that he is dangerous. On April 1, Elmer came to believe a new proposition—the singular proposition about Bugsy

Wabbit that he is not dangerous—but he also steadfastly maintained his belief that Bugsy is dangerous. Thus, on April 1, Elmer believes both that Bugsy is dangerous and that Bugsy is not dangerous. If he reflects on the matter, he may even come to believe the conjunction: that Bugsy Wabbit (the criminal) is dangerous and, in addition, Bugsy Wabbit (my friend) is not dangerous. In fact, if Elmer is sufficiently reflective, he may even *know* that he believes that Bugsy Wabbit is dangerous and, in addition, Bugsy Wabbit is not dangerous. Of course, he will not see any contradiction in this, and not for any cognitive failing on his part.[5]

Kripke objects to this description of the situation on the grounds that "there seem to be insuperable difficulties with [it]: . . . We may suppose that [the subject allegedly having these contradictory beliefs] is a leading philosopher and logician. He would *never* let contradictory beliefs pass. . . . He lacks information, not logical acumen. He cannot be convicted of inconsistency: to do so is incorrect."[6]

On the modified naive theory, our description of Elmer's cognitive state is perfectly acceptable. Part of the reason for this is that, on the modified naive theory, it is misleading to attach logical attributes, such as contradictoriness, to propositions rather than to sentences. 'Hesperus is Hesperus' is a logical truth whereas 'Hesperus is Phosphorus' is not, but the propositions are the same. 'Hesperus is a planet' logically entails 'Hesperus is a planet' whereas 'Phosphorus is a planet' does not, yet the propositions are all the same. 'Hesperus is a planet and Hesperus is not a planet' is a logical contradiction. 'Hesperus is a planet and Phosphorus is not a planet' is logically consistent. The propositions are the same. Logical attributes, such as contradictoriness, apply primarily and in the first instance to sentences in a particular language. (This point is developed further in the following appendix.) Elmer has conflicting, or incompatible, beliefs. One may also say that he has *contradictory* beliefs in the derivative sense that some sentences that encode Elmer's beliefs are contradictory sentences. Of course, even if Elmer is a leading logician he need not realize that his beliefs are contradictory in this derivative sense, for he need not realize that the two sentences he accepts (or utters)—'Bugsy is dangerous' and 'Bugsy is not dangerous'—in fact negate one another. To say that Elmer believes both that Bugsy is dangerous and that he is not dangerous is not, ipso facto, to attribute to Elmer some perverse defect in his reasoning faculty, or anything of the sort. It is only to point out two of his beliefs, without specifying anything about the means by which he grasps the propositions he believes. Elmer is in a state of partial ignorance with respect to Bugsy. There is something important that Elmer does not realize; he is ignorant of some fact.

It might be thought that, for these reasons, Kripke's puzzle is no puzzle on the modified naive theory. David Lewis writes that on the modified naive theory "Kripke's puzzle vanishes," that on that theory "there is no reason to suppose that a leading philosopher and logician would never let contradictory beliefs pass, or that anyone is in principle in a position to note and correct contradictory beliefs if he has them."[7] This response, by itself, is inadequate; as Kripke emphasizes in providing a strengthened version of his puzzle,[8] the case for saying that Elmer does not believe that Bugsy is dangerous seems just as strong as the case for saying that Elmer does believe that Bugsy is dangerous, since Elmer fails to assent to the sentence 'Bugsy Wabbit is dangerous' when he takes it as a sentence concerning his newfound friend and yet he fully and completely understands the sentence. A more complete account of Kripke's puzzle faithful to the spirit of the modified naive theory must invoke the ternary *BEL* relation in order to explain away the temptation to say that Elmer does not believe that Bugsy is dangerous. Elmer withholds belief from the proposition that Bugsy is dangerous, in the sense of 'withhold belief' defined in subsection 8.3.1 above, but strictly speaking he also believes that Bugsy is dangerous.

The strengthened version of Kripke's puzzle relies on a correspondingly strengthened version of the disquotation-principle schema: *A normal English speaker who is not reticent will be disposed to sincere reflective assent to 'S' if and only if he or she believes that S*, where once again the substituends for '*S*' are appropriate English sentences. This quotation-disquotation principle is essentially the biconditional formed from the original disquotation principle and its converse. I have argued that the disquotation principle is unobjectionable, since it is a corollary of the analysis of belief as inward assent. The case is different, however, with respect to its converse, Kripke's quotation principle.[9] A person may believe a proposition by mentally assenting to it when he or she takes it a certain way, and may nevertheless refrain from sincere, reflective outward assent to a standard sentence which expresses that proposition and which he or she fully and completely understands, even though he or she is not reticent and wants very much to reveal his or her beliefs. The person may understand the sentence always by taking the relevant proposition in some other way, not recognizing it as the same proposition that he or she inwardly assents to when it is taken in the first way. (This is precisely the situation in Kripke's original example involving Pierre and the sentence 'London is pretty'.) Failure to give verbal assent is a sign of withheld belief, in the sense used here, but it is not necessarily a sign of failure to believe. In fact, the strengthened version of Kripke's puzzle is a reductio ad absurdum of the quotation principle on which it depends.

Appendix B
Analyticity and A Priority

B.1 Analyticity

It was argued in the preceding appendix that logical attributes such as logical validity (logical truth), consistency, contradictoriness, and entailment apply, primarily and in the first instance, to sentences or sets of sentences (in some cases, ordered sets of sentences), and apply secondarily or derivatively to the propositions and sets of propositions encoded by these sentences. A proposition p may be said to be *logically valid with respect to* a context c and a time t, in the derivative sense, if p is the information content, with respect to c and t, of a logically valid sentence of some possible language. A proposition p may be said to be *logically valid* (*simpliciter*) if it is logically valid with respect to every possible context and every time, i.e., if for every possible context c and time t there is some logically valid sentence of some possible language whose information content with respect to c and t is p.[1]

By contrast, epistemological properties such as a priority, a posteriority, and informativeness apply, primarily and in the first instance, to propositions, or pieces of cognitive information—the objects of knowledge and belief—and apply derivatively to the sentences that encode these propositions. A sentence S may be said to be *a priori with respect to* a context c and a time t, in the derivative sense, if its information content with respect to c and t is a priori in the primary sense, that is, if the information content is in principle knowable solely on the basis of reflection on the concepts (or other proposition components) involved, without recourse to sensory experience. A sentence S may be said to be *a priori* (*simpliciter*) if it is a priori with respect to every possible context and time, i.e., if its information content with respect to any possible context and any time whatsoever is a priori.

The availability of these derivative senses naturally invites such questions as whether all or only logically valid sentences are a priori, and whether all or only logically valid propositions are a priori. One traditional view held that these questions are all to be answered affirmatively. Indeed, it seems quite likely that all logically valid prop-

ositions are a priori, since they should be knowable by reason alone, even if there may be logically valid sentences that are not a priori. (See note 1.) With the failure of the classical program of logicism, it is now generally acknowledged that, in addition to the logical validities, the provable sentences of mathematics are also a priori even though they are not in general logically valid. Thus, the class of a priori sentences includes at least the logical validities and the theorems of mathematics.

The derivative concept of a logically valid proposition presents a certain anomaly that gives rise to yet a third type of a priori sentence: one that is not itself logically valid but whose information content is. This arises because there are sentences that are not themselves logically valid but which have the very same program as a sentence that is. For example, the sentence 'All attorneys are lawyers' has the logical form 'All *F*'s are *G*'s', and hence is not logically valid, yet we may presume it encodes, with respect to any context and time, the very same proposition as the sentence 'All lawyers are lawyers', which is logically valid. Such a sentence—one that is not itself logically valid but whose program coincides with that of a logically valid sentence—shares all the philosophically significant (derivative) epistemological characteristics of a logically valid sentence; it is a priori, certain, knowable by reason alone, and so on. Hence, such sentences might naturally be called upon to perform philosophically significant functions thought to be performed by the logically valid sentences themselves—in epistemological foundationalism, empiricism, reductionism, and so forth. Such sentences fully deserve to be grouped together, for philosophical purposes, with the logically valid sentences. Indeed, traditional philosophy has grouped these two sorts of sentences together, under the rubric 'analytic'. I submit that the *analytic* sentences of a language, as the term is traditionally understood, are precisely those sentences whose programs coincide with that of a logically valid sentence (of some possible language).[2] Indeed, the most conspicuous semantic feature shared by such sentences as 'Attorneys are lawyers' and 'Lawyers are lawyers' is their program. As Quine has pointed out, barring nonextensional devices like quotation marks, a sentence is analytic, in the traditional sense, if it is a logically true sentence or "can be turned into a logical truth by putting synonyms for synonyms" (i.e., if a logically valid sentence can be obtained from it by substitution of a contained term with another of the same meaning).[3] Since the "meaning" (i.e., the program) of such a sentence is unaffected by such a substitution, if the result of such a substitution is logically valid, then the original sentence has the same program as a logically valid sentence. In fact, on certain trivial assumptions concerning the language in question, a sentence is analytic if *and only if*

it "can be turned into a logical truth by putting synonyms for synonyms." A true sentence is *synthetic* if and only if it is not analytic.

The traditional conception of the analytic-synthetic distinction holds that sentences like 'Bachelors are unmarried' and 'A vixen is a female fox' are analytic, and therefore a priori, since they can be turned into logically true sentences by putting 'unmarried man eligible for marriage' for 'bachelor' and 'female fox' for 'vixen', whereas sentences like 'Some bachelors are unhappy' are synthetic. These categorizations, however, must not be taken as definitional of the traditional concepts of analyticity and syntheticity. Rather, they reflect the traditional view that 'vixen' is synonymous with 'female fox' and that 'bachelor' is synonymous with 'unmarried man eligible for marriage' but not with any phrase essentially involving the term 'unhappy'. A different theory concerning the meanings of these or other expressions may yield different verdicts concerning the logico-semantic status of these or other sentences—different categorizations of particular sentences as analytic rather than synthetic or synthetic rather than analytic—without thereby altering the traditional concept of what makes for analyticity.

The modified naive theory is just such a theory. Interestingly, the modified naive theory clashes with traditional categorizations in both ways: Some sentences traditionally considered analytic are counted synthetic, and other sentences traditionally considered synthetic are counted analytic.

It is not surprising that there are sentences traditionally characterized as analytic but regarded as synthetic on the modified naive theory. Two examples are 'Hesperus, if it exists, appears in the evening sky' and Phosphorus, if it exists, appears in the morning sky'.[4] What is more interesting is that there are sentences traditionally considered synthetic but counted analytic on the modified naive theory. Two examples are 'If Cicero was an orator, then so was Tully' and 'Hesperus does not weigh any more than Phosphorus'. In fact, on the modified naive theory, the conditional 'Hesperus, if it exists, is Phosphorus' is analytic, since a logical validity is obtained from it by putting 'Venus' for both 'Hesperus' and 'Phosphorus', and hence it shares the same program with a logically valid sentence. Any logical consequence of 'Hesperus, if it exists, is Phosphorus' is likewise analytic, and encodes (with respect to any context and time) a logically valid proposition.

In *Naming and Necessity*, Kripke made the startling claim that such sentences as these are necessary yet a posteriori. I quote at length, making appropriate insertions in brackets:

> [It is] true that someone can use the name 'Cicero' to refer to Cicero and the name 'Tully' to refer to Cicero also, and not know that

[if he exists, then] Cicero is Tully. So it seems that we do not necessarily know *a priori* that an identity statement between names is true. It doesn't follow from this that the statement so expressed is a contingent one if true. . . .[5]

Are there really [possible] circumstances in which Hesperus [would have existed, but] wouldn't have been Phosphorus? . . . Someone goes by and he calls two *different* stars 'Hesperus' and Phosphorus'. . . . But are those circumstances in which Hesperus [exists, but] is not Phosphorus or [would have existed, but] would not have been Phosphorus? It seems to me that they are not. . . .

[We] can certainly say that the name 'Phosphorus' might not have referred to Phosphorus. We can even say that . . . it might have been the case that . . . something else was . . . *called* 'Phosphorus'. But that still is not a case in which Phosphorus [existed but] was not Hesperus. There might be a possible world in which, a possible counterfactual situation in which, 'Hesperus' and 'Phosphorus' weren't names of the things they in fact are names of. . . . But still that's not a case in which Hesperus [existed, but] wasn't Phosphorus. For there couldn't have been such a case, given that Hesperus is Phosphorus. . . .

[In] advance, we are inclined to say, the answer to the question whether [if it exists, then] Hesperus is Phosphorus might have turned out either way. . . . [There is] one sense in which things might turn out either way, in which it's clear that that doesn't imply that the way it finally turns out isn't necessary. For example, the four color theorem might turn out to be true and might turn out to be false. It might turn out either way. It still doesn't mean that the way it turns out is not necessary. Obviously, the 'might' here is purely 'epistemic'—it merely expresses our present state of ignorance, or uncertainty. . . .[6]

Though for all we know in advance, Hesperus [existed, but] wasn't Phosphorus, [the fact that if Hesperus exists, then it is Phosphorus] couldn't have turned out any other way, in a [metaphysical, non-epistemic] sense. . . . [We] can say in advance, that if Hesperus and Phosphorus are one and the same, then in no other possible world can they [exist and] be different. . . . [In] any other possible world it will be true that [if it exists, then] Hesperus is Phosphorus. So two things are true: first, that we do not know *a priori* that [if it exists, then] Hesperus is Phosphorus, and are in no position to find out the answer except empirically. Second, this is so because we could have evidence qualitatively indistinguishable from the

> evidence we [actually] have and determine the reference of the two names by the positions of two planets in the sky, without the planets being the same. . . .[7]
>
> We have concluded that an identity statement between names, when true at all, is necessarily true, even though one may not know it *a priori*. . . . [There are] possible worlds in which . . . some planet other than Hesperus was called 'Hesperus'. But even so, it would not be a situation in which Hesperus itself [existed but] was not Phosphorus.
>
> Some of the problems which bother people in these situations . . . come from . . . a confusion . . . between what we can know *a priori* in advance and what is necessary. Certain statements—and the identity statement is a paradigm of such a statement on my view—if true at all must be necessarily true. One does know *a priori*, by philosophical analysis, that *if* such an identity statement is true it is necessarily true.[8]

I have no disagreement with Kripke concerning the modal status of such sentences as 'Hesperus, if it exists, is Phosphorus', but I sharply disagree concerning their epistemological status. Kripke defends his thesis that such sentences are necessary by warning against a potential confusion of the fact that Hesperus and Phosphorus are identical (which he alleges to be necessary) with the distinct, semantic fact that the names 'Hesperus' and 'Phosphorus' are co-referential (which is assuredly contingent). He seems not to have heeded this warning in considering the separate question of the epistemological status of these sentences. It is indeed knowable only a posteriori that if 'Hesperus' refers to anything and if 'Phosphorus' refers to anything, then they refer to the same thing. This semantic fact about the names 'Hesperus' and 'Phosphorus' is both contingent and a posteriori. However, the separate, nonlinguistic fact that if Hesperus exists then it is Phosphorus is just the fact that if Venus exists then it is it, and this fact (proposition, "thought", piece of cognitive information) is fully knowable, with complete certainty, by reason alone. Indeed, it is a truth of logic (in the derivative sense). In determining the (derivative) epistemological status of any sentence—in determining whether its content is a priori or a posteriori—it is crucial to bear in mind a sharp distinction between the key notion of semantically encoded information and the entirely irrelevant notion of pragmatically imparted information, such as the information that the sentence in question is true. The semantically encoded information may be knowable a priori even when the sentence's pragmatic impartations are knowable only a posteriori. Since sentences

like 'If Hesperus exists, it is Phosphorus' are analytic, in the traditional sense, it is to be expected that they are metaphysically necessary. By the same token, it should be small wonder that they are also a priori.[9]

B.2 Definition

Central to the modified naive theory is the tenet that, by and large, in any natural language the information values of simple (noncompound) expressions are individuals in the case of singular terms and attribues in the case of predicates, connectives, quantifiers, and sentence-forming operators, whereas, by and large, the information value of a compound expression is a complex consisting of the information values of the expression's simple constituents. The 'by and large' clauses are intended to exclude compound expressions involving nonextensional operators, and perhaps those involving compound predicates and common noun phrases. (See note 4.) They also exclude exceptions generated by explicit semantic stipulation, for example the stipulation that a certain apparently simple expression is to have precisely the same information value as a certain compound one. This central tenet yields the consequence that, in any natural language, by and large, the substitution in a sentence of a compound expression for a simple one, or vice versa, results in different information content—even if the interchanged expressions are co-extensional, so that truth value is preserved. In fact, information content is affected even if the simple expression is *defined* by means of the compound expression, unless the definition is explicitly a strict synonymy definition. Since single words are never compound expressions (except in the special and somewhat rare case of an explicit strict-synonymy stipulation), this means that, by and large, what is expressed by means of a single word cannot be expressed by means of a compound expression in its place. This is true even if the single word and the compound expression are in different natural languages. Thus, if the modified naive theory is indeed correct, as I have argued, by and large a single word of one natural language cannot be exactly translated, preserving meaning (i.e., program), by means of a compound expression, whether of the same or a different natural language.

This last consequence may seem implausible. Alonzo Church, employing Langford's translation test in his critique of Rudolf Carnap, Benson Mates, and Hilary Putnam on "identity of belief," asks the reader to suppose, for the sake of argument, that the single word 'fortnight' has the same meaning in English as the phrase 'period of fourteen days'. In proposing to translate a particular English sentence involving the word 'fortnight' into German, he writes:

> As soon as we set out . . . , our attention is drawn to the fact that the German language has no single word which translates the word 'fortnight', and that the literal translation of the word 'fortnight' from English into German is 'Zeitraum von vierzehn Tagen'. . . . Of course we must ask whether the absence of a one-word translation of 'fortnight' is a deficiency of the German language in the sense that there are therefore some things which can be expressed in English but cannot be expressed in German. But it would seem that it can hardly be so regarded—else we should be obliged to call it a deficiency of German also that there is no word to mean a period of fifty-four days and six hours or that the Latin word 'ero' can be translated only by the three-word phrase 'ich werde sein'. Indeed it should rather be said that the word 'fortnight' in English is not a necessity but a dispensable linguistic luxury.[10]

At least a part of the reason that these remarks concerning the word 'fortnight' seem especially plausible is that the word is a unit measurement term defined in terms of another unit measurement term—to wit, the word 'day' used as a term for a specific period of time. By definition, one fortnight = 14 days. Not all measurement terms can be so defined. For example, the word 'day', in its use as a term for a unit of temporal measurement, is defined not in terms of any other measurement term but as the duration, as of some particular date or epoch d, of one complete rotation of the earth on its axis. Similarly, a term for a unit of measurement of spatial length could be defined as the length of a particular standard bar or stick S as of a certain time t. Philosophical legend has it that the term 'meter' was so defined. This is a useful myth.[11] Given such a definition for 'meter', it would be in the spirit of Church's remarks to claim that the word 'meter', in this sense, is a "dispensable linguistic luxury" of English, and that anything that can be expressed in English using the word 'meter', in this sense, can also be expressed using some phrase such as 'the length of stick S at t'', or a translation thereof. But this simply is not so. One piece of information that can be expressed using the defining phrase is a specification of the length of the stick S at t. The sentence 'S is exactly one meter long' encodes, with respect to the time t, a very different proposition from 'S is the same length as S at t''; the former does, and the latter does not, specify the length of S. Indeed, as Kripke has shown, the two propositions determine different modal intensions; the first is false in any possible world in which S has a different length at t, whereas the second remains true there (and in any other world in which S exists), since S might have been slightly shorter than one meter at t

whereas it is a necessary truth that S, if it exists, has the same length as itself at t.[12] Kripke's modal argument thus shows that, even if the word 'meter' is defined by means of the description 'the length of S at t', their information values remain distinct.

The modal argument could also be made (although perhaps somewhat less forcefully) with respect to 'day', in its use as a term for a unit of temporal length, and 'period of one complete rotation of the earth as of d'. It might have taken exactly $1\frac{1}{8}$ days (i.e., 27 hours) for the earth to complete one rotation on d, though it would remain true that one complete rotation on d takes exactly as much time as it takes. The modal argument cannot be made, however, with respect to Church's original example, since a fortnight could not have been other than exactly fourteen days. This is not because 'fortnight' and 'fourteen days' have the same information value, as Church asks his readers to suppose. They do not;[13] the information value of the phrase 'fourteen days' is a complex made up in part of the number, fourteen, and the length of time, twenty-four hours, whereas the information value of the simple term 'fortnight' is just the length of time, fourteen days.

Kaplan gives the following illuminating argument:

> The yacht owner's guest who is reported by Russell to have become entangled in "I thought that your yacht was longer than it is" should have said, "Look, let's call the length of your yacht a 'russell'. What I was trying to say is that I thought that your yacht was longer than a russell." If the result of such a dubbing [definition] were the introduction of 'russell' as a mere [synonym] for 'the length of your yacht', the whole performance would have been in vain.[14]

The guest stipulated only what the term 'russell' was to refer to, and said nothing whatever about the information value of the term. What, then, is its information value? It is natural to suppose, at least initially, that if a simple singular term is defined by means of a definite description, then, in the absence of any explicit pronouncement concerning information value, the defined term takes on the information value of the description, or perhaps that it has no determinate information value.[15] But this example belies both suppositions. It would seem to be a rule of natural languages that, in the absence of explicit stipulation one way or the other, a simple (noncompound) singular term automatically takes its referent as information value. As with 'russell', so with 'fortnight'. Pace Church, it is indeed a deficiency of German (and English) that there is no word to mean the length of time, fifty-four days and six hours—if it is a deficiency of a language that there is information it does not express. But it is not a serious deficiency, and

if the need arises it is a deficiency easily corrected, as the myth of 'meter' and Kaplan's example of 'russell' amply illustrate. ("Let's call the length of time, fifty-four days and six hours, a 'churchnight'.")

Kaplan's example also suggests something of interest concerning the epistemological status of (the contents of) sentences like 'Your yacht is exactly one russell long'.[16]

Kripke, in addition to his claim that true identity sentences involving distinct proper names are necessary yet a posteriori, made the further claim—even more startling than the first—that there are nontrivial examples of sentences (the contents of) which are true only contingently and yet knowable a priori.[17] Kripke's alleged examples of the contingent a priori arise precisely from situations in which a proper name, or some other simple term, is introduced through stipulating that it is to refer to a particular described object or individual. One of Kripke's examples is the sentence 'Stick *S*, if it exists, is exactly one meter long at time *t*'. According to Kripke, for somone who explicitly stipulates "The word 'meter' is to refer (rigidly) to the length of a particular stick *S* at *t*," the sentence in question is a priori upon the stipulation since the term 'meter' is, in effect, defined so that the sentence would be true. At the same time, the stick *S* might have had a slightly different length at *t*, so that (the content of) the sentence is contingent.

Here again, I agree entirely with Kripke's assessment of the modal status of this sentence, but I disagree sharply concerning the epistemological status. And here again, I would caution against confusing semantically encoded information—the information about the particular length that the stick *S* is just that long at *t*—from pragmatically imparted but not semantically encoded information, such as the information that the sentence 'Stick *S*, if it exists, is exactly one meter long at *t*' is true. The latter information may indeed be a priori for the agent in Kripke's example, and this apparent fact seems to be the source of Kripke's contention that the (content of) the sentence is itself a priori. But even if the former information is a priori for the agent, it does not follow that the latter, semantically encoded information is. In fact, it is not. Others have objected to Kripke's alleged examples of the contingent a priori for these or related reasons.[18] What I wish to emphasize is that the agent's failure to know a priori, solely on the basis of the definition, that *S* is one meter long at *t* is not due to his or her being in no position to know anything (a priori or not) directly about (i.e., *of*, or *de re*) the length, one meter. Indeed, the agent knows a great many things concerning this length: that it is greater than some other particular length (say, two inches), that it is shorter than a mile, and so on. Although the agent is in a very good position to know things directly concerning what is in fact the length of *S*, and although the definition by itself

may result in new (perhaps even a priori) knowledge concerning the word 'meter' and its special relation to the stick *S*, the definition by itself does not and could not result in additional knowledge concerning the very length itself. In particular, the definition does not yield the knowledge concerning the length, 39.3701 inches or one meter, that *S*, if it exists, is exactly *that* long at *t*. To have such knowledge is to know exactly how long *S* is at *t* provided it exists. If the agent in Kripke's example knows the precise length of *S* at *t* on the condition that it exists (I do not say that he or she does; I say *if* he or she does), then surely he or she does not know this a priori, on the basis of the definition, but only a posteriori. Perhaps he or she is in no position to know this at all, not even a posteriori. (Suppose the agent has never seen the stick *S*, and makes the stipulation under the misimpression that the stick is only a fraction of an inch long, or that it is several miles long.) The exact length of *S* at *t*, the fact that it is exactly *that* long (39.3701 inches, or one meter), does not seem to be the sort of thing that could be known a priori. No matter what linguistic maneuvers one performs, no matter what semantics one decrees, it would seem that one cannot know the precise length of *S*—the fact that it is exactly *this* long—simply by reflection on the concepts or object involved, in the way that we know simple arithmetic truths, nor even by reflection on one's stipulations concerning the word 'meter'. Knowledge concerning a particular length that a certain stick (if it exists) is exactly that long would seem to be the paradigm of a posteriori knowledge. It is knowledge gained ultimately by measurement. It would seem that if that is not a posteriori for everyone, then nothing is a posteriori for anyone.[19]

Appendix C
Propositional Semantics

Primitive Vocabulary of the Language $\mathscr{L}$

Punctuation symbols: left and right parentheses, the comma
(Individual) variables: x, y, z, x_1, y_1, z_1, x_2, . . .
Individual constants: a, b, I
First-order monadic predicates: *Bald, Human*
First-order dyadic predicates: *Loves, Identical*
Truth-functional connectives: $\sim$, $\supset$, $\wedge$, $\vee$, $\equiv$
Second-order predicates (quantifiers): $\forall$, $\exists$
Definite description operator: ι
Propositional operators: $\Box$, *Actually, that*
Propositional predicates: *Necessary, Believes*
Temporal operator: *Sometimes*

The well-formed expressions (wfe) of $\mathscr{L}$ are of eight mutually exclusive kinds: singular terms, (first-order) predicates, truth-functional connectives, quantifiers, operators, propositional terms, propositional predicates, and formulas.

Formation Rules of $\mathscr{L}$

1. Any primitive individual constant or variable is a *singular term.*
2. Any primitive first-order monadic predicate is a (*first-order*) *monadic predicate.*
3. If Π is any *monadic predicate* and α is any *singular term,* then $\ulcorner\Pi(\alpha)\urcorner$ is a *formula.*
4. If Π is any primitive dyadic predicate and α and β are any *singular terms,* then $\ulcorner\Pi(\alpha,\beta)\urcorner$ is a *formula.*
5. If ϕ is any *formula,* then so is $\ulcorner\sim\phi\urcorner$.
6. If ϕ and ψ are any *formulas,* then so is $\ulcorner(\phi\supset\psi)\urcorner$.
7. If ϕ and ψ are any *formulas,* then so is $\ulcorner(\phi\wedge\psi)\urcorner$.

8. If ϕ and ψ are any *formulas*, then so is $\ulcorner(\phi \vee \psi)\urcorner$.
9. If ϕ and ψ are any *formulas*, then so is $\ulcorner(\phi \equiv \psi)\urcorner$.
10. If α is any variable and ϕ is any *formula*, then $\ulcorner\alpha\phi\urcorner$ is a *monadic predicate*.
11. If Π is any *monadic predicate*, then $\ulcorner\forall\Pi\urcorner$ is a *formula*.
12. If Π is any *monadic predicate*, then $\ulcorner\exists\Pi\urcorner$ is a *formula*.
13. If Π is any *monadic predicate*, then $\ulcorner\iota\Pi\urcorner$ is a *singular term*.
14. If ϕ is any *formula*, then so is $\ulcorner\Box\phi\urcorner$.
15. If ϕ is any *formula*, then so is $\ulcorner Actually\phi\urcorner$.
16. If ϕ is any *formula*, then $\ulcorner that\phi\urcorner$ is a *propositional term*.
17. If α is any *propositional term*, then $\ulcorner Necessary(\alpha)\urcorner$ is a formula.
18. If α is any *singular term* and β is any *propositional term*, then $\ulcorner Believes\ (\alpha,\beta)\urcorner$ is a *formula*.
19. If ϕ is any formula, then so is $\ulcorner Sometimes\phi\urcorner$.

Semantics for $\mathscr{L}$

Definition of *information value with respect to* a context c and a time t *and under* an assignment of values to variables A in $\mathscr{L}$, $Val_{c,\,t,\,A}$, and of *information-value base with respect to* a context c *and under* an assignment of values to variables A in $\mathscr{L}$, $Valb_{c,\,A}$:

For every context c, time t, and assignment of values to variables A,

1. If α is an individual variable, $Valb_{c,\,A}(\alpha) = Val_{c,\,t,\,A}(\alpha) = A(\alpha)$;
2. $Valb_{c,\,A}(\text{'}a\text{'}) = Val_{c,\,t,\,A}(\text{'}a\text{'})$ = Smith;
3. $Valb_{c,\,A}(\text{'}b\text{'}) = Val_{c,\,t,\,A}(\text{'}b\text{'})$ = Jones;
4. $Valb_{c,\,A}(\text{'}I\text{'}) = Val_{c,\,t,\,A}(\text{'}I\text{'})$ = the agent of c;
5. $Valb_{c,\,A}(\text{'}Bald\text{'})$ = the property of being bald;
6. $Valb_{c,\,A}(\text{'}Human\text{'})$ = the property of being human;
7. $Valb_{c,\,A}(\text{'}Loves\text{'})$ = the relation of loving;
8. $Valb_{c,\,A}(\text{'}Identical\text{'})$ = the relation of identity;
9. $Valb_{c,\,A}(\text{'}\sim\text{'}) = Val_{c,\,t,\,A}(\text{'}\sim\text{'})$ = the property of being the truth value, falsehood;
10. $Valb_{c,\,A}(\text{'}\supset\text{'}) = Val_{c,\,t,\,A}(\text{'}\supset\text{'})$ = the relation *COND:* if u is the truth value truth, then so is v, i.e., a relation which truth bears only to itself and which everything else bears to everything;
11. $Valb_{c,\,A}(\text{'}\wedge\text{'}) = Val_{c,\,t,\,A}(\text{'}\wedge\text{'})$ = the relation of joint truth: $u = v$ = truth;
12. $Valb_{c,\,A}(\text{'}\vee\text{'}) = Val_{c,\,t,\,A}(\text{'}\vee\text{'})$ = the relation of alternative truth: Either u is truth or v is;

13. $Valb_{c,A}(\text{'}\equiv\text{'}) = Val_{c,t,A}(\text{'}\equiv\text{'})$ = the identity relation restricted to truth values, i.e. the relation of being the same truth value;

14. $Valb_{c,A}(\text{'}\forall\text{'})$ = the property *UNIV* of being the universal domain of individuals;

15. $Valb_{c,A}(\text{'}\exists\text{'})$ = the property *EXIST* of being a nonempty set of individuals;

16. $Valb_{c,A}(\text{'}\iota\text{'})$ = the operation O_i of assigning to any singleton set of individuals its only element;

17. $Valb_{c,A}(\text{'}\Box\text{'}) = Val_{c,t,A}(\text{'}\Box\text{'}) = Valb_{c,A}(\text{'}Necessary\text{'}) = Val_{c,t,A}(\text{'}Necessary\text{'})$ = the property of being a necessary truth, i.e., a proposition true in every accessible world;

18. $Valb_{c,A}(\text{'}Actually\text{'}) = Val_{c,t,A}(\text{'}Actually\text{'})$ = the property of being true in w, where w is the possible world in which c occurs;

19. $Valb_{c,A}(\text{'}that\text{'}) = Val_{c,t,A}(\text{'}that\text{'})$ = the operation O_p of assigning every proposition to itself;

20. $Valb_{c,A}(\text{'}Believes\text{'})$ = the relation of believing;

21. $Valb_{c,A}(\text{'}Sometimes\text{'}) = Val_{c,t,A}(\text{'}Sometimes\text{'})$ = the property $\Sigma times$ of obtaining at some time or other.

22. If Γ is a monadic predicate, a dyadic predicate, a quantifier, the definite description operator, or *'Believes'*, then
$Val_{c,t,A}(\Gamma) = \langle Valb_{c,A}(\Gamma), t\rangle$;

23. If Π is a monadic predicate and α is an individual constant or variable, then
$Valb_{c,A}(\ulcorner\Pi(\alpha)\urcorner) = \langle Valb_{c,A}(\alpha), Valb_{c,A}(\Pi)\rangle$
and
$Val_{c,t,A}(\ulcorner\Pi(\alpha)\urcorner) = \langle Val_{c,A,I}(\alpha), Val_{c,t,A}(\Pi)\rangle$;

24. If Π is a dyadic predicate and each of α and β is an individual constant or variable, then
$Valb_{c,A}(\ulcorner\Pi(\alpha,\beta)\urcorner) = \langle Valb_{c,A}(\alpha), Valb_{c,A}(\beta), Valb_{c,A}(\Pi)\rangle$
and
$Val_{c,t,A}(\ulcorner\Pi(\alpha,\beta)\urcorner) = \langle Val_{c,t,A}(\alpha), Val_{c,t,A}(\beta), Val_{c,t,A}(\Pi)\rangle$;

25. If ϕ is any formula, then
$Valb_{c,A}(\ulcorner\sim\phi\urcorner) = Valb_{c,A}(\phi)^\frown\langle Valb_{c,A}(\text{'}\sim\text{'})\rangle$
and
$Val_{c,t,A}(\ulcorner\sim\phi\urcorner) = Val_{c,t,A}(\phi)^\frown\langle Val_{c,t,A}(\text{'}\sim\text{'})\rangle$;

26. If ϕ and ψ are any formulas and C is '$\supset$', '$\wedge$', '$\vee$' or '$\equiv$', then
$Valb_{c,A}(\ulcorner(\phi\ C\ \psi)\urcorner) = Valb_{c,A}(\phi)^\frown Valb_{c,A}(\psi)^\frown\langle Valb_{c,A}(C)\rangle$
and
$Val_{c,t,A}(\ulcorner\phi\ C\ \psi\urcorner) = Val_{c,t,A}(\phi)^\frown Val_{c,t,A}(\psi)^\frown\langle Val_{c,t,A}(C)\rangle$;

27. If Q is either '$\forall$' or '$\exists$' and Π is any monadic predicate, then
$Valb_{c,A}(\ulcorner Q\Pi\urcorner) = \langle Valb_{c,A}(\Pi), Valb_{c,A}(Q)\rangle$

and
$Val_{c,\ t,\ A}(\ulcorner Q\Pi\urcorner) = \langle Val_{c,\ t,\ A}(\Pi),\ Val_{c,\ t,\ A}(Q)\rangle$;

28. If α is any variable and ϕ is any formula, then
$Valb_{c,\ A}(\ulcorner\alpha\phi\urcorner)$ = (the property corresponding to) $\lambda i\ Valb_{c,\ A^{\alpha}_{i}}(\phi)$, where A^{α}_{i} is the assignment that assigns the individual i to α and is otherwise the same as A;

29. If Π is any predicate, then
$Valb_{c,\ A}(\ulcorner\iota\Pi\urcorner) = \langle Valb_{c,\ A}(\text{'}\iota\text{'}),\ Valb_{c,\ a}(\Pi)\rangle$
and
$Val_{c,\ t,\ A}(\ulcorner\iota\Pi\urcorner) = \langle Val_{c,\ t,\ A}(\text{'}\iota\text{'}),\ Val_{c,\ t,\ A}(\Pi)\rangle$;

30. If O is either '$\Box$' or 'Actually' and ϕ is any formula, then
$Valb_{c,\ A}(\ulcorner O\phi\urcorner) = \langle Valb_{c,\ A}(\phi),\ Valb_{c,\ A}(O)\rangle$
and
$Val_{c,\ t,\ A}(\ulcorner O\phi\urcorner) = \langle Val_{c,\ t,\ A}(\phi),\ Val_{c,\ t,\ A}(O)\rangle$
[contrast rule 25 above];

31. If ϕ is any formula, then
$Valb_{c,\ A}(\ulcorner that\phi\urcorner) = \langle Valb_{c,\ A}(\text{'}that\text{'}),\ Valb_{c,\ A}(\phi)\rangle$
and
$Val_{c,\ t,\ A}(\ulcorner that\phi\urcorner) = \langle Val_{c,\ t,\ A}(\text{'}that\text{'}),\ Val_{c,\ t,\ A}(\phi)\rangle$;

32. If Γ is a monadic predicate and α is a compound singular term (one obtained by means of formation rule 13), or Γ is 'Necessary' and α is a propositional term, then
$Valb_{c,\ A}(\ulcorner\Gamma(\alpha)\urcorner) = Valb_{c,\ A}(\alpha)^{\frown}\langle Valb_{c,\ A}(\Gamma)\rangle$
and
$Val_{c,\ t,\ A}(\ulcorner\Gamma(\alpha)\urcorner) = Val_{c,\ t,\ A}(\alpha)^{\frown}\langle Val_{c,\ t,\ A}(\Gamma)\rangle$;

33. If Π is a dyadic predicate, α is a simple singular term, and β is a compound singular term, then
$Valb_{c,\ A}(\ulcorner\Pi(\alpha,\beta)\urcorner) = \langle Valb_{c,\ A}(\alpha)\rangle^{\frown}Valb_{c,\ A}(\beta)^{\frown}\langle Valb_{c,\ A}(\Pi)\rangle$;
$Val_{c,\ t,\ A}(\ulcorner\Pi(\alpha,\beta)\urcorner) = \langle Val_{c,\ t,\ A}(\alpha)\rangle\ Val_{c,\ t,\ A}(\beta)^{\frown}\langle\ Val_{c,\ t,\ A}(\Pi)\rangle$;
$Valb_{c,\ A}(\ulcorner\Pi(\beta,\alpha)\urcorner) = Valb_{c,\ A}(\beta)^{\frown}\langle Valb_{c,\ A}(\alpha),\ Valb_{c,\ A}(\Pi)\rangle$;
and
$Val_{c,\ t,\ A}(\ulcorner\Pi(\beta,\alpha)\urcorner) = Val_{c,\ t,\ A}(\beta)^{\frown}\langle Val_{c,\ t,\ A}(\alpha),\ Val_{c,\ t,\ A}(\Pi)\rangle$;

34. If Π is a dyadic predicate and α and β are both compound singular terms, then
$Valb_{c,\ A}(\ulcorner\Pi(\alpha,\beta)\urcorner) = Valb_{c,\ A}(\alpha)^{\frown}Valb_{c,\ A}(\beta)^{\frown}\langle Valb_{c,\ A}(\Pi)\rangle$
and
$Val_{c,\ t,\ A}(\ulcorner\Pi(\alpha,\beta)\urcorner) = Val_{c,\ t,\ A}(\alpha)^{\frown}Val_{c,\ t,\ A}(\beta)^{\frown}\langle Val_{c,\ t,\ A}(\Pi)\rangle$;

35. If α is an individual constant or variable and β is a propositional term, then
$Valb_{c,\ A}(\ulcorner Believes(\alpha,\beta)\urcorner)$
$= \langle Valb_{c,\ A}(\alpha)\rangle^{\frown}Valb_{c,\ A}(\beta)^{\frown}\langle Valb_{c,\ A}(\text{'}Believes\text{'})\rangle$

and
$Val_{c,\ t,\ A}(\ulcorner Believes(\alpha,\beta)\urcorner)$
$= \langle Val_{c,\ t,\ A}(\alpha)\rangle \frown Val_{c,\ t,\ A}(\beta) \frown \langle Val_{c,\ t,\ A}(\text{'Believes'})\rangle$;

36. If α is a compound singular term and β is a propositional term, then
$Valb_{c,\ A}(\ulcorner Believes(\alpha,\beta)\urcorner) = Valb_{c,\ A}(\alpha) \frown Valb_{c,\ A}(\beta) \frown \langle Valb_{c,\ A}(\text{'Believes'})\rangle$
and
$Val_{c,\ t,\ A}(\ulcorner Believes(\alpha,\beta)\urcorner)$
$= Val_{c,\ t,\ A}(\alpha) \frown Val_{c,\ t,\ A}(\beta) \frown \langle Val_{c,\ t,\ A}(\text{'Believes'})\rangle$;

37. If ϕ is any formula, then
$Valb_{c,\ A}(\ulcorner Sometimes\phi\urcorner) = \langle Valb_{c,\ A}(\phi), Valb_{c,\ A}(\text{'Sometimes'})\rangle$
and
$Val_{c,\ t,\ A}(\ulcorner Sometimes\phi\urcorner) = \langle Valb_{c,\ A}(\phi), Val_{c,\ t,\ A}(\text{'Sometimes'})\rangle$.
[Note that the information value of $\ulcorner Sometimes\phi\urcorner$ is made up, in part, of the information value base of ϕ rather than its information value.]

Definition: For every information value p, expression e of $\mathscr{L}$, context c, time t, and assignment of values variables A, p is the *information content* of *e with respect to c and t under A* iff e is a formula of $\mathscr{L}$ and $p = Val_{c,\ t,\ A}(e)$.

Definition: For every information value base p, expression e of $\mathscr{L}$, context c, and assignment of values to variables A, p is the (*information*) *content base* of *e with respect to c under A* iff e is a formula of $\mathscr{L}$ and $p = Valb_{c,\ A}(e)$.

Definition of *information value base with respect to* a context c (*simpliciter*) in $\mathscr{L}$, $Valb_c$, and of *information value with respect to* a context c *and* a time t (*simpliciter*) in $\mathscr{L}$, $Val_{c,\ t}$: For every wfe e of $\mathscr{L}$, and every object o, $Valb_c(e) = o$ iff for every assignment of values to variables A, $Valb_{c,\ A}(e) = o$, and $Val_{c,\ t}(e) = o$ iff, for every assignment of values to variables A, $Val_{c,\ t,\ A}(e) = o$.

Definition of *information value with respect to* a context c (*simpliciter*) in $\mathscr{L}$, Val_c: For every wfe e of $\mathscr{L}$, and every object o, $Val_c(e) = o$ iff $Val_{c,\ \text{the time of}\ c}(e) = o$.

Definition: For every wfe e of $\mathscr{L}$, and every assignment of values to variables A, the *program* of e *under* $A = \lambda c\ Valb_{c,\ A}(e)$.

Definition: For every wfe e of $\mathscr{L}$, and every function f, the *program* of e (*simpliciter*) $= f$ iff, for every assignment of values to variables A, the program of e under $A = f$.

Definition of *the eternalization with respect to* a time t *of* a content base p, of $\mathscr{L}$, $Etern_t(p)$: For every content base p, content p', and

time t, $Etern_t(p) = p'$ iff, for every formula ϕ of $\mathscr{L}$ and context c, $Valb_c(\phi) = p$ iff $Val_{c,t}(\phi) = p'$.

Definition: A *p-structure* is any 10-tuple, $\mathfrak{A} = \langle U, W, w_@, R, T, D, \mathscr{I}, \omega, \tau, S\rangle$, where

(i) U is a nonempty set of individuals [the set of possible individuals of $\mathfrak{A}$];

(ii) W is a nonempty set of worlds;

(iii) $w_@ \in W$ [the actual world of $\mathfrak{A}$];

(iv) R is a reflexive and symmetric relation on W [accessibility in $\mathfrak{A}$];

(v) T is a nonempty set of times;

(vi) $D \in U^{W \times T}$ [the assignment of domains to world-time pairs of $\mathfrak{A}$];

(vii) $\mathscr{I} \in U\{\text{all individuals}\}$;

($viii$) $\omega \in W\{\text{all worlds}\}$;

(ix) $\tau \in T\{\text{all times}\}$;

and

(x) S is a function whose domain is {the property of being bald, the property of being human, the relation of loving, the relation of believing} and such that

(a) if P is a property, $S(P) \in (U^W)^T$ and, for all $t \in T$ and $w \in W$, $S(P)(t)(w) \subseteq D(w,t)$;

(b) S (loving) $\in ((U \times U)^W)^T$ and for all $t \in T$, $w \in W$, $S(\text{loving})(t)(w) \subseteq D(w,t) \times D(w,t)$;

(c) S (believing) $\in ((U \times \{\text{information contents of formulas of } \mathscr{L}\})^W)^T$ and for all $t \in T$, $w \in W$, $S(\text{believing})(t)(w) \subseteq D(w,t) \times \{\text{information contents of formulas of } \mathscr{L}\}$.

Definition of property and relation *exemplification with respect to a* time t *and* a world w *in* a p-structure $\mathfrak{A}$, $exemp^{\mathfrak{A}}_{t,w}$ and $exemp^{\mathfrak{A}}_{w}$, and of *truth with respect to* a time t *and* a world w *in* a p-structure $\mathfrak{A}$, $true^{\mathfrak{A}}_{t,w}$ and $true^{\mathfrak{A}}_{w}$, for information values and information value bases of $\mathscr{L}$: For every *p-structure* $\mathfrak{A} = \langle U, W, w_@, R, T, D, \mathscr{I}, \omega, \tau, S\rangle$, individuals $i, j \in U$, time $t \in T$, and world $w \in W$,

1. If P is either the property of being bald or the property of being human, then

i *exemp*$^{\mathfrak{A}}_{t,w}$ P iff $i \in S(P)(t)(w)$, and, for every time t', i *exemp*$^{\mathfrak{A}}_{w}$ $\langle P,t'\rangle$ iff $i \in S(P)(\tau(t'))(w)$.

2. $\langle i,j\rangle$ *exemp*$^{\mathfrak{A}}_{t,w}$ loving iff $\langle i,j\rangle \in S(\text{loving})(t)(w)$, and, for every time t', $\langle i,j\rangle$ *exemp*$^{\mathfrak{A}}_{t,w}$ $\langle$loving, $t'\rangle$ iff $\langle i,j\rangle \in S(\text{loving})(\tau(t'))(w)$.

3. $\langle i,j\rangle$ *exemp*$^{\mathfrak{A}}_{t,w}$ identity iff $i = j$, and, for every time t', $\langle i,j\rangle$ *exemp*$^{\mathfrak{A}}_{w}$ $\langle$identity, $t'\rangle$ iff $i = j$.

4. If i' is any individual and P is any property, then $\langle i',P\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff $\mathscr{I}(i')$ *exemp*$^{\mathfrak{A}}_{t,w}$ P, and, for every time t', $\langle i',\langle P,t'\rangle\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff $\mathscr{I}(i')$ *exemp*$^{\mathfrak{A}}_{w}$ $\langle P,t'\rangle$.

5. If i' and j' are any individuals and P is either loving or identity, then $\langle i',j',P\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff $\langle \mathscr{I}(i'),\mathscr{I}(j')\rangle$ *exemp*$^{\mathfrak{A}}_{t,w}$ P, and, for every time t', $\langle i',j',\langle P,t'\rangle\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff $\langle \mathscr{I}(i'), \mathscr{I}(j')\rangle$ *exemp*$^{\mathfrak{A}}_{w}$ $\langle P, t'\rangle$.

6. If p is any content base, then $p \frown \langle$being falsehood$\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff p is not *true*$^{\mathfrak{A}}_{t,w}$. If p is any information content, then $p \frown \langle$being falsehood$\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff p is not *true*$^{\mathfrak{A}}_{w}$.

7. If p and q are any content bases, then $p \frown q \frown \langle COND\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff, if p is *true*$^{\mathfrak{A}}_{t,w}$, then q is *true*$^{\mathfrak{A}}_{t,w}$. If p and q are any information contents, then $p \frown q \frown \langle COND\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff, if p is *true*$^{\mathfrak{A}}_{w}$, then q is *true*$^{\mathfrak{A}}_{w}$.

8. If p and q are any content bases, then $p \frown q \frown \langle$joint truth$\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff both p and q are *true*$^{\mathfrak{A}}_{t,w}$. If p and q are any information contents, then $p \frown q \frown \langle$joint truth$\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff both p and q are *true*$^{\mathfrak{A}}_{w}$.

9. If p and q are any content bases, then $p \frown q \frown \langle$alternative truth$\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff either p is *true*$^{\mathfrak{A}}_{t,w}$ or q is *true*$^{\mathfrak{A}}_{t,w}$. If p and q are any information contents, then $p \frown q \frown \langle$alternative truth$\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff either p is *true*$^{\mathfrak{A}}_{w}$ or q is *true*$^{\mathfrak{A}}_{w}$.

10. If p and q are any content bases, then $p \frown q \frown \langle$identity$\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff, p is *true*$^{\mathfrak{A}}_{t,w}$ iff q is *true*$^{\mathfrak{A}}_{t,w}$. If p and q are any information contents, then $p \frown q \frown \langle$identity$\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff, p is *true*$^{\mathfrak{A}}_{w}$ iff q is *true*$^{\mathfrak{A}}_{w}$.

11. If P is any property, then $\langle P, UNIV\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff every element of $D(w, t)$ *exemp*$^{\mathfrak{A}}_{t,w}$ P. If P is any property and t' and t'' are any times, then $\langle\langle P, t'\rangle,\langle UNIV, t''\rangle\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff every element of $D(w,t'')$ *exemp*$^{\mathfrak{A}}_{w}$ $\langle P,t'\rangle$.

12. If P is any property, then $\langle P, EXIST\rangle$ is *true*$^{\mathfrak{A}}_{t,w}$ iff at least one element of $D(w, t)$ *exemp*$^{\mathfrak{A}}_{t,w}$ P. If P is any property and t' and t'' are any times, then $\langle\langle P,t'\rangle ,\langle EXIST,t''\rangle\rangle$ is *true*$^{\mathfrak{A}}_{w}$ iff at least one element of $D(w,t'')$ *exemp*$_{w}$ $\langle P,t'\rangle$.

13. If f is any singulary function from individuals to information content bases of $\mathscr{L}$, then i *exemp*$^{\mathfrak{A}}_{t,w}$ (the property corresponding to)

f iff $f(i)$ is $true^{\mathfrak{U}}_{t,\,w}$, and, for every time t, i $exemp^{\mathfrak{U}}_{w}$ ⟨(the property corresponding to) f, t⟩ iff $f(i)$ is $true^{\mathfrak{U}}_{\tau(t),\,w}$.

14. If P and P' are any properties, then $\langle O_i, P, P'\rangle$ is $true^{\mathfrak{U}}_{t,\,w}$ iff exactly one element of $D(w, t)$ $exemp^{\mathfrak{U}}_{t,\,w}$ P and it also $exemp^{\mathfrak{U}}_{t,\,w}$ P'. If P and P' are any properties and t_1, t_2, and t_3 are any times, then $\langle\langle O_i,t_1\rangle,\langle P, t_2\rangle,\langle P',t_3\rangle\rangle$ is $true^{\mathfrak{U}}_{w}$ iff exactly one element of $D(w, \tau(t_1))$ $exemp^{\mathfrak{U}}_{w}$ $\langle P, \tau(t_2)\rangle$ and it $exemp^{\mathfrak{U}}_{w}$ $\langle P', \tau(t_3)\rangle$.

15. If i' is any individual, P is any property, and P' is either loving or identity, then $\langle i',O_i,P,P'\rangle$ is $true^{\mathfrak{U}}_{t,\,w}$ iff exactly one element j' of $D(w, t)$ $exemp^{\mathfrak{U}}_{t,\,w}$ P and $\langle \mathcal{I}(i'),j'\rangle$ $exemp^{\mathfrak{U}}_{t,\,w}$ P'; for any times t_1, t_2, and t_3, $\langle i',\langle O_i,t_1\rangle,\langle P,t_2\rangle\rangle,\langle P',t_3\rangle\rangle$ is $true_{w}$ iff exactly one element j' of $D(w, \tau(t))$ $exemp^{\mathfrak{U}}_{w}$ $\langle P, \tau(t_2)\rangle$ and $\langle \mathcal{I}(i'), j'\rangle$ $exemp^{\mathfrak{U}}_{w}$ $\langle P,\tau(t_3)\rangle$; $\langle O_i,P,i',P'\rangle$ is $true^{\mathfrak{U}}_{t,\,w}$ iff exactly one element j' of $D(w,t)$ $exemp^{\mathfrak{U}}_{t,\,w}$ P and $\langle j',\mathcal{I}(i')\rangle$ $exemp^{\mathfrak{U}}_{t,\,w}$ P'; and for any times t_1, t_2, and t_3, $\langle\langle O_i,t_1\rangle,\langle P,t_2\rangle,i',\langle P',t_3\rangle\rangle$ is $true^{\mathfrak{U}}_{w}$ iff exactly one element j' of $D(w, \tau(t_1))$ $exemp^{\mathfrak{U}}_{w}$ $\langle P, \tau(t_2)\rangle$ and $\langle j', \mathcal{I}(i')\rangle$ $exemp^{\mathfrak{U}}_{w}$ $\langle P', \tau(t_3)\rangle$.

16. If P_1 and P_2 are any properties and P_3 is either loving or identity, then $\langle O_i,P_1,O_i,P_2,P_3\rangle$ is $true^{\mathfrak{U}}_{t,\,w}$ iff exactly one element i' of $D(w,t)$ $exemp^{\mathfrak{U}}_{t,\,w}$ P_1, exactly one element j' of $D(w,t)$ $exemp^{\mathfrak{U}}_{t,\,w}$ P_2, and $\langle i',j'\rangle$ $exemp^{\mathfrak{U}}_{t,\,w}$ P_3; and for any times t_1, t_2, t_3, t_4, t_5, $\langle\langle O_i,t_1\rangle, \langle P_1,t_2\rangle, \langle O_i,t_3\rangle, \langle P_2,t_4\rangle, \langle P_3,t_5\rangle\rangle$ is $true^{\mathfrak{U}}_{w}$ iff exactly one element i' of $D(w, \tau(t_1))$ $exemp_{w}$ $\langle P_1, \tau(t_2)\rangle$, exactly one element j' of $D(w, \tau(t_3))$ $exemp_{w}$ $\langle P_2, \tau(t_4)\rangle$, and $\langle i',j'\rangle$ $exemp_{w}$ $\langle P_3, \tau(t_5)\rangle$.

17. If p is any information content or information-content base, then ⟨p, being necessary⟩ is $true^{\mathfrak{U}}_{t,\,w}$ iff p is $true^{\mathfrak{U}}_{t,\,w'}$, for every world w' such that wRw'; ⟨p, being necessary⟩ is $true^{\mathfrak{U}}_{w}$ on exactly the same condition; and ⟨O_p,p,being necessary⟩ is $true^{\mathfrak{U}}_{t,\,w}$ on exactly the same condition, and $true^{\mathfrak{U}}_{w}$ on exactly the same condition.

18. If p is any information content or information-content base and w' is any world, then ⟨p,being true in w'⟩ is $true^{\mathfrak{U}}_{t,\,w}$ iff p is $true^{\mathfrak{U}}_{t,\,\omega(w')}$ and ⟨p,being true in w⟩ is $true^{\mathfrak{U}}_{w}$ on exactly the same condition.

19. If i' is any individual and p is any content base, then ⟨i',O_p,p,believing⟩ is $true^{\mathfrak{U}}_{t,\,w}$ iff
⟨$I(i')$, $Etern_t(p)$⟩ ∈ S(believing)(t)(w).
If i' is any individual, p is any information content, and t' is any time, then ⟨i',O_p,p,⟨believing,t'⟩⟩ is $true^{\mathfrak{U}}_{w}$ iff
⟨$I(i')$,p⟩ ∈ S(believing)($\tau(t')$)(w).

20. If P is any property and p is any content base, then ⟨O_i,P,O_p,p,believing⟩ is $true^{\mathfrak{U}}_{t,\,w}$ iff exactly one element i' of $D(w,t)$ $exemp^{\mathfrak{U}}_{t,\,w}$ P and
⟨i',$Etern_t(p)$⟩ ∈ S(believing)(t)(w).

If P is any property, p is any information content, and t_1, t_2, and t_3 are any times, then $\langle\langle O_i,t_1\rangle,\langle P,t_2\rangle,O_p,p,\langle\text{believing},t_3\rangle\rangle$ is $true^{\mathfrak{A}}_w$ iff exactly one element i' of $D(w, \tau(t_1))$ $exemp^{\mathfrak{A}}_w$ $\langle P, \tau(t_2)\rangle$ and $\langle i',p\rangle \in S$ (believing)$(\tau(t_3))(w)$.

21. If p is any content base, then $\langle p,\Sigma times\rangle$ is $true^{\mathfrak{A}}_{t,\,w}$ iff there is a time $t' \in T$ such that p is $true^{\mathfrak{A}}_{t',\,w}$, and $\langle p,\Sigma times\rangle$ is $true^{\mathfrak{A}}_{w}$ on exactly the same condition.

Definition of *truth in* a p-structure $\mathfrak{A}$, $true^{\mathfrak{A}}$, for information contents of $\mathscr{L}$: For any information content p and p-structure $\mathfrak{A} = \langle U, W, w_@, R, T, D, \mathscr{I}, \omega, \tau, S\rangle$, p is $true^{\mathfrak{A}}$ iff p is $true^{\mathfrak{A}}_{w_@}$.

Definition of *entailment* for information contents of $\mathscr{L}$: For every set K of information contents and every information content p, K *entails* P iff, for every p-structure $\mathfrak{A}$, if every element of K is true$^{\mathfrak{A}}$, then so is p.

Definition of *validity* for information contents of $\mathscr{L}$: For any information content p, p is *valid* iff $\varnothing$ entails p (where $\varnothing$ is the empty set of information contents).

Notes

Introduction

1. This argument is given more or less explicitly by G. E. Moore in "Facts and Propositions," in his *Philosophical Papers* (New York: Collier Books), at pp. 67–68. Moore does not use the terminology of 'singular propositions', nor does he focus primarily on formulations that are unquestionably de re, but he does speak of "general facts" and "non-general facts," and the intent of his argument is clearly in the Fregean spirit of the argument given here.
2. The argument that follows is derived from one advanced by David Kaplan in the preface to his "Demonstratives" (draft number 2; unpublished) for the conclusion that, in the possible-world semantics of quantified modal logic, singular propositions seem to be needed in the analysis of propositions expressed by sentences involving quantification across a modal operator, even if the proposition expressed by such a sentence is not itself singular but general. The general point is made more forceful and more forcefully, however, when it is applied to the locutions of de re propositional attitude. There is much less resistance nowadays (and there was less in the late 1960s and the 1970s) to the use of singular propositions as the modal content of sentences for the purpose of evaluating the modal value of a sentence (as necessary, possible, etc.) than there is (and has been) to the idea of treating singular propositions as the cognitive contents of sentences for the purpose of evaluating the truth value of propositional-attitude attributions. Moreover, the successful development of possible-world semantics allows one the freedom to evaluate and analyze first-order modal sentences without considering the question of the proposition contained in (expressed by) a sentence. A sentence $\ulcorner\Box S\urcorner$ is true, and its component S therefore necessary, if and only if S is true with respect to every accessible world w. If S happens to be of the form $\ulcorner\Pi(a)\urcorner$, where Π is a monadic predicate and a is a singular term, then $\ulcorner\Box S\urcorner$ is true if and only if, for every accessible world w, Π applies with respect to w to the referent of a with respect to w. Nothing need be said directly concerning the content (or even concerning the corresponding intension, i.e., the correlated function from possible worlds to truth values) of S itself. The situation is not analogous in the case of a propositional-attitude attribution $\ulcorner b$ believes that $S\urcorner$, where it seems required to consider the cognitive content of S as a separate semantic value that bears on the truth value of the attribution.
3. A possible Fregean response to this argument will be discussed in section 9.2.

Chapter 1

1. See also Wittgenstein, *Philosophical Grammar* (Berkeley: University of California Press, 1974), at pp. 315–318; *Philosophical Investigations* (New York: Macmillan), at p. 216.

For a more recent endorsement of the general strategy see Ian Hacking, "Comment on Wiggins," in *Philosophy of Logic*, ed. S. Körner (Berkeley: University of California Press, 1976).

2. A word of clarification is needed concerning my use of the semantic predicates 'encode' and 'information'. Throughout this book I am concerned with discrete units of information that are specifiable by means of a 'that'-clause, e.g. the information that Socrates was wise. These discrete units are *pieces of information*. I shall generally use the mass noun 'information' as if it were shorthand for the count noun phrase 'piece of information', i.e., as a general term whose extension is the class of pieces of information. Thus, I write 'information that is such-and-such' to mean "pieces of information that are such-and-such," 'the same information' to mean "the same pieces of information," 'different information to mean "different pieces of information," and so on. I use the verb 'encode' in such a way that an unambiguous declarative sentence encodes (with respect to a given possible context *c*) a *single* piece of information, which is referred to (with respect to *c*) by the result of prefixing 'the information that' to the sentence and which is to be called 'the information content' of the sentence (with respect to *c*). A declarative sentence may encode (with respect to a given context) two or more pieces of information, but if it does so it is ambiguous. Pieces of information encoded by the logical consequences of an unambiguous sentence are not themselves encoded, in this sense, by the sentence. The (piece of) information that snow is white and grass is green is different information (a different piece of information) from the (piece of) information that snow is white, though intuitively the latter is included as part of the former. The sentence 'Snow is white and grass is green' encodes only the former, not the latter. This constitues a departure from at least one standard usage, according to which the information content of a sentence is perhaps something like a class of pieces of information, closed under logical consequence.

   I am not concerned in this book with a notion of an *amount* of information, which arises in the mathematical theory of communication or information. The information *that snow is white and grass is green and Socrates is Socrates* may be no more or less information than the information *that both snow is white if and only if grass is green and either snow is white or grass is green*. Nevertheless, general considerations involving Leibniz's Law strongly suggest that they are numerically distinct pieces of information. For instance, the first concerns Socrates whereas the second does not.

3. I use here, and throughout this book, a quasi-technical notion of the *context* of an utterance, which is such that, for any particular actual utterance of any expression by anyone, if any facts had been different in any way, even if they are only facts entirely independent of and isolated from the utterance itself, then the context of the utterance would, ipso facto, be a different context—even if the utterance is made by the very same speaker in the very same way to the very same audience at the very same time in the very same place. To put it another way, although a single utterance occurs in indefinitely many different possible worlds, any particular possible context of an utterance occurs in one and only one possible world, so that, in every possible world in which the same utterance occurs, it occurs in a new and different context—even if the speaker, his or her manner of uttering, the time of the utterance, the location of the speaker, the audience being addressed, and all other such features and aspects of the utterance remain exactly the same. There is a very good reason for using the term 'context' in this way: Suppose, for example, that it will come to pass that a Democrat is elected to the presidency in the year 2000, and consider a possible world *W* that is exactly like the actual world in every detail up to January 1, 1999, but in which a Republican is elected to the presidency in 2000. Suppose I here and now utter the sentence

Actually, a Republican will be elected to the presidency in A.D. 2000.

In the actual world, I thereby assert a piece of information that is necessarily false. In *W*, on the other hand, I thereby assert a piece of information that is necessarily true. I utter the very same sequence of words of English, with the very same English meanings, in both possible worlds, yet I assert different things. If we were to use the term 'context' in such a way that the context of my utterance remains the same in both worlds, we should be forced to say, quite mysteriously, that the sentence I uttered is such that it would have encoded different information with respect to the context in which I uttered it if *W* had obtained even though both its meaning and its context of utterance would remain exactly the same. Using the term 'context' as I shall throughout this book, we may say instead that, although the very same utterance occurs by me both in *W* and in the actual world, the context of the utterance is different in the two worlds. This allows us to say that the sentence I utter takes on different information contents with respect to *different* contexts of utterance, thereby assimilating this phenomenon to the sort of context sensitivity that is familiar in cases of sentences like 'A Republican is presently president'.

4. The latter clause is needed in order to distinguish 'Bill loves Mary' from 'Mary loves Bill', where the sequential order of composition is crucial. This succinct statement of the rule connecting sentences and their information contents is only an approximation to the truth. A complicated difficulty arises in connection with the latter clause of the rule and with such quantificational locutions as 'someone' in 'Someone is wise'. Grammatically the sentence 'Someone is wise' is analogous to 'Socrates is wise', though logically and semantically they are disanalogous. In 'Socrates is wise', the predicate 'is wise' attaches to the singular term 'Socrates'. This situation is reversed in 'Someone is wise', wherein the restricted quantifier 'someone' attaches to the predicate 'is wise'. Thus, whereas grammatically 'someone' is combined with 'is wise' to form the first sentence in just the same way that 'Socrates' is combined with 'is wise' to form the second sentence, the information values of 'someone' and 'is wise' are combined very differently from the way the information values of 'Socrates' and 'is wise' are combined. (This complication may lie behind at least part of Russell's motivation for claiming that definite descriptions and quantificational locutions generally ("denoting phrases") have no "meaning in isolation." See note 2 to chapter 3.

A perhaps more important qualification to the general rule is noted in the next paragraph of the text. Yet another important qualification concerns overlaid quantifiers. It is necessary to distinguish between the information contents of such constructions as

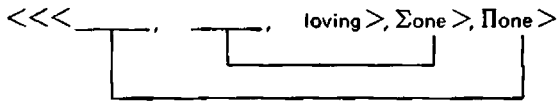
(A) For everyone x there is someone y such that x loves y

and

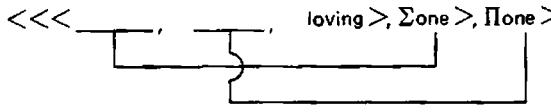
(B) For everyone x there is someone y such that y loves x.

One possible method is to include as part of the structure of some propositions *unoccupied positions* as well as some device that connects or links a proposition component, such as the information value of a quantifier, with an unoccupied position within the proposition. Thus, for example, the information value (with respect to a time t) of the open sentence 'x loves y' may be taken to be something like a partially defined ordered triple, with the first two places undefined and the third place filled by the relation of loving: ⟨_____, _____, loving⟩. (Actually the third place would be filled by the temporally indexed relation of *loving at t*: ⟨_____, _____, ⟨loving, t⟩⟩—see subsection 2.2.5.) Entities of this sort may be called *open propositions*. (They are not to be confused with proposition matrices, which will be defined below. A

proposition matrix may have the form of a totally defined ordered n-tuple, whereas open propositions always have the form of partially defined ordered n-tuples, or sequences containing partially defined ordered n-tuples as elements, etc.) The information values of the quantifiers 'for everyone' and 'there is someone such that' are certain higher-order properties. (Specifically, they are properties of one-place functions from individuals to truth values.) Let us designate them by 'Πone' and 'Σone', respectively. Then, on this method, the information content of sentence A is not merely a sequence of the information values of the information-valued components of the sentence, but something that is also interconnected. This proposition might be diagrammed thus:



The information content of sentence B, on the other hand, would be diagrammed thus:



(Strictly accurate diagrams would replace 'loving' with '⟨loving, t⟩'.) The lines of connection, or links, though not strictly "elements" of the propositions, are regarded as essential to the structures of the propositions. They serve both to *close* what would otherwise be open propositions and to distinguish them from each other. See W. V. O. Quine, *Mathematical Logic* (Cambridge, Mass.: Harvard University Press, 1965), at p. 70; Gareth Evans, "Pronouns, Quantifiers and Relative Clauses (I)," in *Reference, Truth and Reality*, ed. M. Platts (London: Routledge and Kegan Paul, 1980).

On this method, the open sentence 'y loves x' has exactly the same information value as 'x loves y'. Hence, on this method, sentences A and B are such that either can be obtained from the other simply by replacing one component—the contained open sentence with two free variables—by another with the same information value. Nevertheless, A and B clearly differ in information content. This result conflicts with a certain compositionality principle, commonly attributed to Frege, according to which the information value of a compound expression such as A or B is a function solely of the information values of its information-valued components. [See Alonzo Church, "Intensional Isomorphism and Identity of Belief," *Philosophical Studies* 5, no. 5 (1954): 65–73, for a similar but more sharply articulated principle.] On this method, the connections between the information value of a quantifier and the corresponding unoccupied positions within the information value of the open sentence, as indicated by the containing sentence, must also be taken into account.

A superior method, due to Church and pointed out to me by David Kaplan, eliminates open propositions in favor of Russellian propositional functions, i.e., functions from an individual to a singular proposition involving that individual. (Church himself applies the general method in such a way as to invoke only Fregean functions from pure concepts of individuals to Fregean purely general propositions, although the general method can also accommodate anti-Fregean theories by invoking propositional functions.) On this method, the information content of sentence A is regarded as

having the structure of the following ordered pair:

$\langle(\lambda x)\langle(\lambda y)\langle x, y, \text{loving}\rangle, \Sigma\text{one}\rangle, \Pi\text{one}\rangle$.

The information content of sentence B, on the other hand, is regarded as having the following structure:

$(\lambda x)\langle(\lambda y)\langle y, x, \text{loving}\rangle, \Sigma\text{one}\rangle, \Pi\text{one}\rangle$.

(Here again, for complete accuracy, 'loving' should be replaced by '⟨loving, *t*⟩'.) The first element of the former proposition is the propositional function that assigns to each individual *x* the proposition made up of the propositional function that assigns to each individual *y* the proposition that *x* loves *y* and the second-order property Σone. The first element of the latter proposition is the appropriate analogue. A variant of this method (closer to the spirit of the naive and modified naive theories defined below) replaces these propositional functions with the properties of loving someone (at *t*) and of being loved by someone (at *t*), respectively. The information content of sentence A may be regarded as the following complex proposition: that the function that assigns truth to an individual *x* if *x* loves someone or other, and assigns falsehood otherwise, assigns truth to everyone whatsoever. This powerful method need not assign any information value to an open sentence like '*x* loves *y*', except relative to an assignment of values to its free variables, and hence generates no counterexamples to the original compositionality principle. Even on this method, however, it is not true in general that the information value of a compound expression involving bound variables is a complex made up entirely of the information values of its information-valued components. If the first element of the information content of sentence A is the information value of any component of the sentence—for example, the component '*x* there is someone *y* such that *x* loves *y*'—then the information value of that component is not made up of the information values of its information-valued components.

5. See also the introduction above and note 4 to chapter 6 below.

Chapter 2

1. See Alonzo Church, "Review of Carnap's *Introduction to Semantics*," *Philosophical Review* 52 (1943): 298–304, at pp. 299–301; Kurt Gödel, "Russell's Mathematical Logic," in *The Philosophy of Bertrand Russell*, ed. P. A. Schilpp (New York: Tudor, 1944), at pp. 128–129. The particular argument concerning predicates is given in greater detail, for the special case of common nouns, in my *Reference and Essence* (Princeton University Press, 1981), at pp. 48–52.
2. See William Kneale and Martha Kneale, "Propositions and Time," in *G. E. Moore: Essays in Retrospect*, ed. A. Ambrose and M. Lazerowitz (New York: Humanities Press, 1970), at p. 235; Mark Richard, "Temporalism and Eternalism," *Philosophical Studies* 39 (1981): 1–13.
3. The length of the time interval is a vague matter. For many purposes it may be taken to be the entire year of 1971. It should be noted that when the time interval involved in a proposition is significantly long, the proposition may mimic its noneternal matrix—for example, in contexts like 'Mary once believed that Nixon was a Republican, and she still believes that'—as long as one stays within the boundaries of the time interval in question. Relatively stable properties (such as being a Republican, as opposed to being president) tend to lengthen the time interval in question. This is similar to a point made by Kneale and Kneale, "Propositions and Time," at pp. 232–233.
4. On Frege's theory, the domain of this function would consist of senses that determine times, rather than the times themselves.

   There is on Frege's theory no reason why the time indication or time specification

that supplements the incomplete present-tense sentence could not be verbal, e.g. 'At 12:00 noon on July 4, 1983, this tree is covered with green leaves'. This aspect of Frege's theory allows for a solution to the problem of failure of substitutivity of co-referential singular terms in temporal contexts—a solution very different from Frege's solution to the parallel problem of failure of substitutivity in propositional attitude contexts. Consider the following example. The expressions 'the U.S. president' and 'Ronald Reagan' refer to the same individual with respect to the time of my writing these words, but the former cannot be substituted salva veritate for the latter in the true sentence 'In 1978, Ronald Reagan was a Republican'. Substitution yields 'In 1978, the U.S. president was a Republican', which is false on the relevant reading (the Russellian secondary occurrence or small scope reading), since in 1978 Jimmy Carter was president and a Democrat. Frege may solve this problem, not implausibly, by noting that the expression 'the U.S. president' is incomplete and requires supplementation by a time specification, such as may be provided by the time of utterance, before it can refer to an individual. The description 'the U.S. president', supplemented by the time of my writing these words, refers to the same individual as the name 'Ronald Reagan'. Supplemented by the year 1978, or by a verbal specification thereof, it refers to Jimmy Carter. The sentence 'In 1978, the U.S. president was a Republican' includes a verbal time specification, 'in 1978', which we may assume, supersedes the time of utterance in completing any expression occurring within its scope in need of completion by a time specification. This solution is dissimilar from Frege's treatment of substitutivity failure in propositional-attitude contexts. On Frege's theory, a propositional-attitude operator such as 'Jones believes that' creates an oblique context in which expressions refer to their customary information values ("senses") instead of their customary referents. On the Fregean solution to substitutivity failure in temporal contexts presented here, the referent of 'the U.S. president', as occurring within the context 'in 1978 _____', is just its customary referent.

5. For example, the meaning of the term 'table' might include, in addition to its program, some sort of conceptual content, e.g. a specification of the function of a table. If so, it does not follow that this sort of conceptual entity is any part of the information value of the term. Nor does it follow that it is *analytic*, in something like the classical sense, that tables have such-and-such a function. (See appendix B below on analyticity on the modified naive theory.) What does follow is that, in order to know fully the meaning of 'table', one would have to know that the things called 'tables' are conventionally believed to have such-and-such a function.
6. The need for double indexing was apparently first noted in 1967 by Hans Kamp in unpublished material distributed to a graduate seminar while Kamp was a graduate student at UCLA. See his "Formal Properties of 'Now'," *Theoria* 37 (1972): 227–273. Kamp's results were reported in A. N. Prior's " 'Now'," *Noûs* 2 (1968): 101–119.
7. See Mark Richard, "Tense, Propositions, and Meanings," *Philosophical Studies* 41 (1982): 337–351, at pp. 346–349. The idea of doubly indexing information content to both contexts and times is Richard's.
8. The value base of the result of attaching an information-content operator (such as 'necessarily' or the 'that'-operator) to a sentence is a complex made up of the value base of the operator and the content base of the sentence, rather than its information content. Thus, for example, the value base of the 'that'-clause 'that Frege is busy' with respect to any context c does not involve the information content of 'Frege is busy' with respect to c, which is the proposition that Frege is busy at t, where t is the time of c. Instead, it is something like the ordered pair of two elements: (1) A certain abstract entity, analogous to a property which is the operation of assigning any proposition to itself (this operation—call it 'O_p'—is the value base of the 'that'-

operator with respect to any context); and (2) the proposition matrix *fb*. Thus, the value base of 'that Frege is busy' has the following structure: $\langle O_{p}, \langle$Frege, being busy$\rangle\rangle$. The information content of 'Sometimes Frege believes that he is busy' has the following structure, where 'Σtimes' designates the property of proposition matrices of being true at some time(s):

$\langle\langle$Frege, O_{p}, $\langle$Frege, being busy$\rangle$, believing$\rangle$, Σtimes$\rangle$.

A more detailed account is provided in appendix C. For one possible account of verb tenses in accordance with the modified naive theory, as developed here, see my "Tense and Information Content," in *Propositions and Attitudes*, ed. N. Salmon and S. Soames (forthcoming).

Chapter 3

1. See James Cappio, "Russell's Philosophical Development," *Synthèse* 46 (1981): 185–205.
2. Russell claimed that quantifier phrases ("denoting phrases") such as 'some man', unlike genuine singular terms, have no "meaning in isolation," i.e., no information value. It may seem that this claim is due to an oversight on Russell's part, since, for any such quantificational construction, Russell's (largely implicit) semantics for higher-order logic includes clearly identifiable higher-order propositional functions to serve as the contribution of the construction to the information encoded by the sentences in which it figures—for example, the propositional function *f* that is the information value of '$(\lambda F)(\exists x)[x$ is a man & $F(x)]$', i.e., the function *f* that assigns to any singularly propositional function *F* of individuals the proposition that something both is a man and instantiates *F*. See Richard Montague's "The Proper Treatment of Quantification in Ordinary English," in his *Formal Philosophy*, ed. R. Thomason (New Haven: Yale University Press, 1974), and appendix A of Kaplan's "Opacity," in *The Philosophy of W. V. O. Quine*, ed. P. A. Schilpp (forthcoming). Russell was initially impressed with the fact that a restricted quantificational construction, e.g. 'some man', may serve as the grammatical subject of a sentence, e.g. 'Some man is wise'. (See note 4 to chapter 1.) He may have feared that if 'some man' was assigned a "meaning in isolation," then the sentence 'Some man is wise' would have to be construed on the model of 'Plato is wise', ascribing the property of wisdom to the "meaning in isolation" of the grammatical subject. It is possible that Russell, in denying that restricted quantificational constructions have "meaning," meant only to emphasize that, though a description (whether definite, like 'the author of *The Republic*', or indefinite, like 'a pupil of Socrates') is a noun phrase that may serve as the grammatical subject of a sentence and which is replaceable by a name like 'Plato', its logical form—as opposed to its grammatical form—is not that of a singular term to which the first-order predicate attaches but rather that of a second-order predicate that attaches to the first-order predicate. Since Russell separates logic from grammar, however, this observation offers insufficient grounds for depriving restricted quantificational constructions of information value.

There is a better reason for Russell to have denied that such quantifier phrases as 'some man', 'all men', and 'some unique author of *The Republic*' have information value. On Russell's theory, the phrase 'some man' is construed as having the logical form of a compound second-order predicate constructed from the existential quantifier 'for something', an individual variable, and an open compound second-order predicate: 'for something *x*, *x* is a man and *x* _____', or, more formal, '$(\exists x)[x$ is a man & x _]'. Indeed, Russell called quantifier phrases 'incomplete symbols'. A complete sentence like 'Some man is wise' is construed as having the logical form of 'For something *x*, *x* is a man and *x* is wise'. Assuming that Russell's theory employs Church's method

(cited in note 4 to chapter 1) for dealing with quantifiers, the existential quantifier 'for something' is regarded as a simple second-order predicate, while the remainder of the sentence '*x*, *x* is a man and *x* is wise' is regarded as a closed compound first-order predicate, synonymous with 'is both a man and wise' (roughly synonymous with 'is a wise man'). The sentence 'Some man is wise' would thus be regarded as having the logical form of 'Something is both a man and wise', whose information content has the structure of the ordered pair of the first-order propositional function that is the information value of 'is both a man and wise' and the second-order propositional function that is the information value of 'something'. There is no place in this proposition for the propositional function *f* mentioned above. Hence, there is nothing which the phrase 'some man' contributes on its own to the proposition that some man is wise (i.e., the proposition that something is both a man and wise), although the information value of 'man' figures indirectly in the construction of the information value of 'is both a man and wise'. (The theory of open propositions suggested in note 4 to chapter 1 provides an information value for "incomplete" quantifier phrases like '$(\exists x)[x$ is a man & x _____]', but the information content this theory provides for '$(\exists x)[x$ is a man & x is wise]' is very different from that provided by Russell's theory.) Interpreting Russell's Theory of Descriptions in this way may better reveal the deeper import of Russell's claim that 'the author of *The Republic*' has no "meaning in isolation" and, in particular, does not have Plato as its meaning. The central tenet of the Theory of Descriptions is that a definite description such as 'the author of *The Republic*' is semantically equivalent to its corresponding uniqueness-restricted indefinite description 'some unique author of *The Republic*'. The *meaning* of a sentence-component, in Russell's usage, is the proposition-component which the sentence-component contributes to propositions, i.e., its information value. Given Russell's adherence to a variant of the singly modified naive theory, only a genuine name of an individual *x* has *x* as its meaning in this sense. A definite description for *x* is only an "incomplete symbol," and thus has no meaning at all, not even *x* itself, in this sense. For a more detailed discussion of these matters, see my "Reference and Information Content: Names and Descriptions," in *Handbook of Philosophical Logic IV*, ed. D. Gabbay and F. Guenthner (Dordrecht: Reidel, forthcoming).

From the viewpoint of English syntax, neither Russell's theory of 'some' nor the proposal concerning the propositional function *f* mentioned above seems decisively superior to the simple theory that 'some' is a dyadic quantifier whose information value is a binary relation between one-place functions from individuals to truth values.

3. David Kaplan exposes and exploits some of these important points of contact in "How to Russell a Frege-Church," *Journal of Philosophy* 72, no. 19 (1975): 716–729.
4. In characterizing the sense of an expression as a purely conceptual entity, I intend the term 'concept' with a more or less ordinary meaning and not with that of Frege's special use of 'Begriff'. Senses are neither empirically observed (as are external, concrete objects) nor "had" in the way that sensations or other private experiences are had, but are abstract entities that are "grasped" or "apprehended" by the mind. In addition, I intend the term 'pure' to exclude concepts that include nonconceptual elements as constituents. A genuine sense may involve reference to an object, but it must do so by including a conceptual representation of the object in place of the object itself. This will be clarified in subsection 5.1.1.
5. The correspondence is not an exact coincidence. The Fregean thought that Socrates is (timelessly) wise is supposed to be made up of a purely conceptual mode of presentation of Socrates (i.e., the sense of 'Socrates') together with the sense of 'is wise', which is a function from modes of individual presentations to Fregean thoughts,

namely, to the thought whose "component parts" are the mode of individual presentation and the function itself. Every Fregean thought consists partly of a function from senses to thoughts. (On this, see the second to last paragraph of Frege's "Über Begriff und Gegenstand.") Roughly, in lieu of these functions from Fregean senses, Russellian propositions employ *propositional functions*, i.e., functions from objects (not necessarily senses or any other kind of intensional entity) to singular propositions (or perhaps proposition matrices)—namely, to the singular proposition (matrix) whose "constituents" are the object and the propositional function itself. Russellian propositions may be regarded as replacing even Fregean individual concepts (singular term senses) with complex propositional functions. (See note 2 above.) Thus, even purely general Russellian propositions are built from entities—propositional functions—that play no role in the construction of a Fregean thought. (See the exchange between Frege and Russell concerning the nature of propositions or "thoughts" reproduced in Frege, *Philosophical and Mathematical Correspondence*, ed. Gabriel et al. (University of Chicago Press, 1980), at pp. 149–170. See especially pp. 163 and 169–170 on the nature of the "thought" that Mont Blanc is more than 4,000 meters high and "what is actually asserted" by the sentence 'Mont Blanc is more than 4,000 meters high'.)

More recent philosophy has fused Russellian propositional functions and functions from concepts to propositions to form a new proposition-building device that is more versatile than either of these and similar in many respects to the "class concepts" of Russell's *Principles of Mathematics*. A predicate like 'is wise' is taken to contribute something that combines with purely conceptual individual concepts to form general propositions and with objects (or with individual concepts having objects as constituents) to form singular propositions. My own doctrine, which builds on the modified naive theory, replaces Fregean individual concepts (singular term senses) with complexes consisting of two parts: (1) an abstract entity, analogous to a property, which is the operation O_1 of assigning to any one-place function from individuals to truth values the unique individual to which the function assigns truth, if there is one, and assigning nothing otherwise (this operation is the information value of 'the' in the singular); and (2) the information value of a common noun or of a common noun phrase (generally, a temporally indexed property or kind). These complexes are the information values of definite descriptions. They serve many of the purposes for which Fregean individual concepts were introduced, but they are constructed solely from elements of the modified naive theory. The information values of definite descriptions are unlike Frege's individual concepts (singular term senses) in that they frequently have individuals as constituents—to wit, whenever the constitutive property does.

6. Linsky's more recent book *Oblique Contexts* (University of Chicago Press, 1983) contains similar remarks:

> Frege introduced his distinction between sense and reference in the context of an attempt to answer the questions 'How can a true statement of identity ever be informative?' and 'What information does it convey?' [In the standard sort of case] it is entirely clear what has been learned when [the subject] learns that Hesperus is Phosphorus. What he has learned is that the criterion of identification associated with 'Hesperus', its sense, picks out the same object as is picked out by the criterion of identification associated with 'Phosphorus', *its* sense. (p. 132)
>
> What we discover when we learn that Hesperus = Phosphorus is that the alternative ways we use for fixing the reference of these names picks out the same planet. (p. 146)

7. The question of what it is on Fregean theory for a thought to be "about" an individual can be a large one, depending on the special sense of 'about' that is in question. One such question on which Frege was largely tacit is the question of what conditions a

Fregean thought must satisfy in order for its belief to constitute a de re belief about an individual—the topic of an extensive literature. The project David Kaplan undertakes in "Quantifying In," in *Reference and Modality*, ed. L. Linsky (New York: Oxford University Press, 1971), may be seen as an attempt to answer this question. From the point of view of the singly modified naive theory, which Kaplan now favors, the main project undertaken in "Quantifying In" was ill conceived. A Fregean thought is never a de re (i.e., singular) proposition.

8. The direct-reference theory is discussed in subsection 5.1.1 below.
9. S. Kripke, *Naming and Necessity* (Cambridge, Mass.: Harvard University Press, 1982), pp. 107–108; B. Russell and A. N. Whitehead, *Principia Mathematica* (Cambridge University Press, 1963), introduction, chapter III, p. 67; David Wiggins, "Frege's Problem of the Morning Star and the Evening Star," in *Studies on Frege II: Logic and the Philosophy of Language*, ed. M. Schirn (Stuttgart: Bad Canstatt, 1976).
10. Mill held a complex theory of information value, according to which the information encoded by a sentence like 'Socrates is wise' has at least two components: (1) the proposition about Socrates that he has the property of wisdom, and (2) the metalinguistic proposition about the expressions 'Socrates' and 'wise' that the individual referred to ("denoted") by the former has the property "connoted" by the latter and is therefore among the things "denoted" by the latter. This would mean that any term will itself form part of its own information value. In the special case of an identity sentence like 'Hesperus is Phosphorus', Mill held that the first component is null, so that the information encoded reduces to the metalinguistic truth that the name 'Hesperus' "denotes" the same thing as the name 'Phosphorus'. Thus, Mill solved Frege's Puzzle in a manner similar to that of Frege's early theory. Mill went beyond the early Frege, though, generalizing the metalinguistic solution to Frege's Puzzle in such a way as to make a similar solution available for the other puzzles that arise on the naive theory. It is doubtful that this was part of Mill's motivation in propounding his dual component theory of information value, though it doubtless would have been seen as lending independent support to the theory. By the same token, many of the objections to Frege's early view apply a fortiori to the metalinguistic component of Mill's theory of information value.
11. See chapter 1 of my *Reference and Essence* (Princeton University Press, 1981), where this claim is developed in greater detail.
12. This is not the most generous interpretation possible, though I find it the most plausible. One might take Frege's argument instead as a reductio ad absurdum of the metalinguistic analysis of identity in conjunction with the naive theory. But there is virtually nothing in his discussion of the metalinguistic analysis in "Über Sinn und Bedeutung" to indicate that this was Frege's intent, and it is unlike Frege to omit such matters of detail when doing so might engender misunderstanding. I cannot help thinking that, if Frege had consciously thought that the position that he was arguing against was *wrong* in supposing that the matter of which thing a name names is entirely the result of human decision or usage, he would have said so, and his argument would have proceeded quite differently.

    Kaplan has suggested reconstructing Frege's criticism of the metalinguistic analysis so that it is based on the weaker claim that the circumstance of two names naming the same thing is due in part, though not entirely, to an accident of human culture (whereas the fact that Hesperus is Phosphorus is entirely independent of human culture). This weaker claim is compatible with Frege's theory of sense. In fact, I have argued that it is a consequence of Frege's theory of sense. Here again, taken as a piece of Fregean scholarship (though I believe this is not Kaplan's intent), this more generous interpretation is implausible. Frege says that the fact that two names are

co-referential is arbitrary, and that a sentence conveying this fact would not concern the object so named but would concern only the names. Moreover, this interpretation robs Frege of his insight that the circumstances of a name naming some particular individual is constituted entirely by some sort of linguistic activity (e.g. semantic stipulation, convention, usage) and not at all by the individual's being one of a certain conceptually indicated kind, such as being a heavenly body visible in such-and-such location at such-and-such time.

13. The stipulational, de re character of the naming of an individual does not mean that the responsibility may be taken lightly. Quite the contrary. See T. S. Eliot's "The Naming of Cats," in his *Old Possum's Book of Practical Cats* (New York: Harcourt, Brace, 1939), pp. 1–2.

Chapter 4

1. R. Sharvy, "Why a Class Can't Change Its Members," *Noûs* 2, no. 4 (1968): 303–314.
2. One pair of sentences proposed to me as a counterinstance to Frege's Law in correspondence by a prominent philosopher of logic and semantics is ' 'Hesperus' refers to Hesperus' (uninformative) and ' 'Hesperus' refers to Phosphorus' (informative). I shall criticize the claim that Frege's Law *could be* (let alone that it *is*) subject to counterexample, but perhaps a special caveat is called for in connection with this particular example. These two sentences are equally informative, in the sense of the term 'informative' that is relevant to Frege's Puzzle. In particular, even the first sentence is informative—and not simply because it entails the nontrivial fact that 'Hesperus' is not nonreferring. The sentence 'If 'Hesperus' refers to anything, it refers to Hesperus' is equally informative, in the relevant sense. (In this connection, see Kripke, *Naming and Necessity*, pp. 68–70.) These sentences are not only informative, they (or, more accurate, their information contents) are the subject of serious dispute among semanticists. Richard Montague denied that 'Hesperus' refers to Hesperus, as did Russell. (See note 2 to chapter 3, above.) If a semantic oracle were to have pronounced the truth that 'Hesperus' does indeed refer to Hesperus, these philosophers should have found the pronouncement only too painfully informative. (More probably, they might have denounced the oracle as a fraud.) The sentence ' 'Hesperus' refers to Hesperus' is, by itself and in abstraction from context, incomplete. Reference is a relation among expressions, objects, and linguistic systems; names refer to things (or fail to refer to things) in this or that language, or in this or that idiolect. There are (possible) languages in which 'Hesperus' refers to nothing, and still others in which it refers to the Milky Way. The information that 'Hesperus' refers to Hesperus in English is a nontrivial piece of information about English. Things might have been otherwise, and it is not "given" or known a priori what the expression 'Hesperus' refers to in English.

   One sentence that might be correctly regarded as uninformative, in the relevant sense, and is easily confused with ' 'Hesperus' refers to Hesperus', is the following: 'The sentence ' 'Hesperus' refers to Hesperus in English' is true in English*', where 'English*' refers to the extension of English into a metalanguage for English. The apparent triviality of this meta-metatheoretic sentence is no doubt the source of the erroneous claim that ' 'Hesperus' refers to Hesperus' is uninformative. But, as we shall see, it is crucial in discussing Frege's Puzzle to maintain a sharp distinction between the information content of a sentence *S* and the further and separate metalinguistic information that *S* is true. Frege's Puzzle concerns the former, and not generally the latter. The reasons behind the apparent triviality of the meta-metatheoretic sentence mentioned above are complex (see Kripke, *Naming and Necessity*, pp. 68–70),

but in no way does this sentence present a problem for Frege's Law. That the two original sentences are equally informative does not entail that they semantically encode the same piece of information. It does entail that the modified-naive-theorist, in claiming that they encode the same information, has nothing to fear from Frege's Law.

3. R. Carnap, "Reply to Leonard Linsky," *Philosophy of Science* 16, no. 4 (1949): 347–350, at pp. 347–348.

4. One way of attempting to block the argument for this conclusion is to deny that the information encoded by 'Hesperus is Hesperus' employs the same mode of composition of its elements as the information encoded by 'Hesperus is Phosphorus'. Such a denial is strongly suggested by the proposal, made some time ago by Hilary Putnam, that the information encoded by a sentence is a function both of the information values of the sentence components and of the very logical structure of the sentence itself. See Putnam, "Synonymity, and the Analysis of Belief Sentences," *Analysis* (April 1954): 114–122, at pp. 118–119, and especially note 8. Very recently, David Kaplan has revived this idea by claiming (in conversation and in lecture) that, where a and b are co-referential simple singular terms, the information content of $\ulcorner a = a \urcorner$ includes some indication, by way of something like the *lines of connection* or *links* introduced in note 4 to chapter 2 above, of the recurrence of the same individual, whereas the information content of $\ulcorner a = b \urcorner$ includes only the recurrence with no separate indication of the recurrence.

This proposal has the somewhat implausible consequence that, concerning any individual, there is information that can be expressed using two (or more) completely synonymous names for that individual, but simply cannot be expressed at all using only one term (fewer) for that individual, even if the one term is completely synonymous with both of the two original terms. It would follow that any language that is just like English except for including one additional name completely synonymous with (one of) the existing English name(s) for a given individual is, ipso facto, expressively richer than English; though the name seems a superfluous addition to the language, instead English is expressively deficient in that it lacks some such extra name. [See Alonzo Church, "Intensional Isomorphism and Identity of Belief," *Philosophical Studies* 5, no. 5 (1954): 65–73, at pp. 70–71.] Ironically, the problem with the weaker language would be that it expresses too much structure rather than too little. In fact, this proposal imputes more structure to the proposition that Hesperus is Hesperus than there is in the sentence 'Hesperus is Hesperus', since the sentence includes two occurrences of the same proper name for the same individual but apparently includes no separate indication, by way of lines of connection or the like, of this recurrence. Even the Fregean fine-grained theory of information content does not impute this much extra structure to the information contents of sentences of the form $\ulcorner a = a \urcorner$. Some argument is needed to make it plausible that the information content of $\ulcorner a = a \urcorner$ should have this extra structure that does not correspond to any additional structure in the sentence itself.

Moreover, this argument cannot rely on the sorts of considerations that generate Frege's Puzzle, for the puzzle reappears on this proposal through consideration of pairs of sentences like 'Hesperus appears in the evening sky' (uninformative) and 'Phosphorus appears in the evening sky' (informative), where there is no term recurrence. In fact, since there is apparently no separate indication in a sentence of the form $\ulcorner a = a \urcorner$ that the same term a for the same individual is occurring twice, the proposal does not completely remove the puzzle even where there is term recurrence. To consider an example due to Kripke ["A Puzzle About Belief," in *Meaning and Use*, ed. A. Margalit (Dordrecht: Reidel, 1979), at pp. 265–266], suppose that Peter er-

roneously believes that the name 'Paderewski' as used to refer to the famous pianist and the name 'Paderewski' as used to refer to the Polish nationalist leader and statesman are not the same name for the same individual, but instead the same syntactic sound and shape being used to refer to two distinct individuals: the musician and the statesman. For Peter, the sentence 'Paderewski is Paderewski' is informative if uttered while pointing to two pictures of Paderewski, one in his musician guise and and the other in his statesman guise, whereas the same sentence is uninformative if uttered while pointing to the same picture twice. The additional structure Kaplan imputes to the proposition does not account for the difference in informativeness, since even on this proposal the very same proposition is expressed on both occasions. (This sort of consideration will be discussed at greater length in chapter 5.)

Chapter 5

1. The parenthetical phrase 'semantically correlated individual' in this statement of the common ground of Frege and Russell respects Russell's theoretical claim that ordinary proper names do not have any "meaning in isolation." See note 2 to chapter 3. The higher-order propositional functions mentioned there could serve as the conceptual representations mentioned here. Even on Russell's theory, there is some semantic relation, albeit not reference or the relation he called 'meaning' (i.e., information value), between a surface English expression like 'Socrates', in an ordinary use, and the man, Socrates. This is the relation that correlates with each uniqueness-restricted existential quantifer ⌈some unique ϕ⌉ the unique object that satisfies ϕ, if there is one, and nothing otherwise. Russell misleadingly called this relation 'denotation'. This semantic relation might be seen as the relative product of two more basic relations: a semantic relation between the expression and a certain higher-order propositional function, which may serve as a conceptual representation, and a further nonsemantic relation of fit between this representation and the relevant object. In this respect, Russell's notion of "denotation" for definite descriptions may be structurally analogous to Frege's conception of reference, or Bedeutung.
2. Frege and Russell wrote before the advent of modern intensional semantics, and consequently neither spoke of reference or truth with respect to a possible world or with respect to a time, but only of reference ("meaning") or truth (in a language) simpliciter. The parenthetical phrase 'with respect to w and t' indicates the natural and usual extensions of their account to modal and temporal semantics. However, both Frege and Russell treated the phenomenon of tense and other temporal operators differently from the usual treatment of today, and neither clearly distinguished tense from the distinct phenomenon of indexicality. See note 6 to chapter 2.
3. See my *Reference and Essence*, pp. 14–21, 43–44, and 54–55, for a more detailed discussion of these notions.
4. See Keith Donnellan, "Proper Names and Identifying Descriptions," in *Semantics of Natural Language*, ed. D. Davidson and G. Harman (Dordrecht: Reidel, 1972); Donnellan, "Speaking of Nothing," in *Naming, Necessity, and Natural Kinds*, ed. S. Schwartz (Ithaca, N.Y.: Cornell University Press, 1977); Kripke, *Naming and Necessity*, especially pp. 5–15, 20–21, 24–34, 48–49, 53–97, 127–128, 134–135, and 160–163; David Kaplan, "Demonstratives," draft no. 2 (unpublished ms., Dept. of Philosophy, UCLA, 1977); Hilary Putnam, "The Meaning of 'Meaning'," in his *Mind, Language and Reality: Philosophical Papers* 2 (Cambridge University Press, 1975).
5. H. Putnam, "Meaning and Reference," *Journal of Philosophy* 70 (November 8, 1973): 699–711, at pp. 700–704.
6. See pp. 219–222 of Putnam's "The Meaning of 'Meaning' " on his notion of a "purely psychological state." The relevant notion of a purely conceptual or qualitative concept

is discussed in my *Reference and Essence* on pp. 19–20 and 54–55. See note 4 to chapter 3 of the present book.

7. Hilary Putnam, in his "Comments" on Kripke's "Puzzle About Belief" in *Meaning and Use*, ed. A. Margalit (Dordrecht: Reidel, 1979), writes (at p. 285): "99% of the sameness of meaning is sameness of reference. Certainly Frege's [Puzzle] shows meaning cannot just *be* reference, but there may be much more truth than falsity to the view that meaning is reference." On Putnam's view, the meaning of a term is typically composed partly (mostly) of its extension (really its intension—the corresponding function from possible worlds to extensions; see my *Reference and Essence*, p. 153, note 41), and partly (though apparently only, at most, 1%) of its associated *stereotype*, which is a kind of conceptual content.
8. See Michael Devitt, *Designation* (New York: Columbia University Press, 1981).
9. See Hilary Putnam, "Meaning and Reference," *Journal of Philosophy* 70 (November 8, 1973): 699–711, at pp. 704–706, and note 7 above.
10. This proposal is Kaplanesque in spirit. In his "Demonstratives," Kaplan, using Fregean terminology, suggests identifying the "cognitive significance" (a common translation of Frege's 'Erkenntniswerte') of an indexical with its character, and he speaks of the character as a "manner of presentation" of the information value ("content"). Nevertheless, the proposal is not endorsed by Kaplan, who, as a direct-reference theorist sympathetic to something like the singly modified naive theory, argues that the information value ("proposition component") of a single-word indexical singular term is best identified with its referent.
11. See Saul Kripke, "A Puzzle About Belief," in *Meaning and Use*, ed. A. Margalit (Dordrecht: Reidel, 1979), at pp. 263–265.
12. There is a certain difficulty in applying the Generalized Frege Strategy against the Fregean theory, owing to vagueness in what may be included in the "purely conceptual content" of a term. It may be objected, for instance, that my concept of an elm tree includes the concept of being called 'elm' in English, and perhaps even the concept of being a different genus from the things called 'beeches' in English. If so, my concepts of elms and beeches are significantly different after all. Strictly speaking, these semantic concepts are not purely conceptual, since quotation names like ' 'elm' ' and ' 'beeches' ', and the term 'English', are nondescriptional. The concept of being called 'elm' in English involves both the term 'elm' and the English language as constituents, and hence it is not the sort of concept that Fregean theory posits as the information values of terms.

    Perhaps there is a purely conceptual analogue to the concept of being called 'elm' in English, one that replaces the term 'elm' and the English language by purely conceptual representations. If there is such a concept, is it part of the purely conceptual content I attach to the term 'elm'? In the relevant sense of 'conceptual content', it is not. Not everything one believes about elms can be part of the information value of the term 'elm', or of the conceptual or representational content attached to the term 'elm', as the notion of conceptual or representational content is intended in Fregean theory. Otherwise, every sentence *S* sincerely uttered by someone involving the word 'elm' would be such that the conditional ⌜If there are any elms, then *S*⌝ is analytically true. One could not acquire new beliefs expressed by means of the term 'elm', and hence one could not change one's mind about anything expressed in terms of 'elm' (e.g. that Jones is standing by an elm tree) without literally changing the subject. Hence, at least some of the things believed about elms are not part of the information value of the term 'elm'. In particular, there are compelling reasons for denying that any concept like that of being called such-and-such in English can be a part of the information value of terms like 'elm' and 'beech'. It is no truth of logic, for example,

that elms are called 'elms' in English. See note 2 to chapter 4. The objections to the verbal theory of information value also apply here. Whatever the information value of 'elm' is, there are other terms (e.g. in other languages) that have the same information value. For example, the information that Jones is standing by an elm tree can be expressed in German using either the word 'Ulme' or the word 'Rüster' for elm. Surely the information value of 'Rüster' in German does not include any concept of what things of that kind are called in English. A German speaker may know what an elm is—may have a concept of an elm tree—without having the foggiest idea what elms are called in English. Also, for most terms (e.g. 'tree', 'clock', 'teacher', 'acrobat', 'dentist'), it is distinctly implausible to suppose that the information value of the term includes the concept of being so-called in English. Each is perfectly translatable into any number of languages. The typical German speaker knows what a tree or a clock is—has the concept of a tree or a clock—even if he or she does not know or have any opinion as to the English term for a tree or a clock. There is no reason why 'elm' should be different from 'tree' in this respect. See Kripke, *Naming and Necessity*, pp. 68–70; "A Puzzle About Belief," p. 274, note 12.

Chapter 6

1. In claiming that Frege and Russell and their followers have mistaken pragmatically imparted information for semantically encoded information, I do not mean that they would assent to such things as 'The sentence 'Hesperus is Phosphorus' expresses in English the information about itself that it is true'. Clearly they would not; in any case, they need not. Nor would someone who mistakes a particular celebrity impersonator for the president of the United States assent to 'The president is the celebrity impersonator'. Philosophers mistake pragmatically imparted information for semantically encoded information in failing to keep the two sharply distinct and consequently judging whether a sentence *S* is informative partly on the basis of information pragmatically imparted by utterances of *S*.

   Other writers have drawn distinctions similar to the one drawn here between semantically encoded and pragmatically imparted information as part of a defense of something like the original or the modified naive theory, though I came upon the idea independently. See Michael Tye, "The Puzzle of Hesperus and Phosphorus," *Australasian Journal of Philosophy* 56, no. 3 (1978): 219–224, at p. 224; Raymond Bradley and Norman Swartz, *Possible Worlds: An Introduction to Logic and Its Philosophy* (Indianapolis: Hackett, 1979), at pp. 191–192; Tom McKay, "On Proper Names in Belief Ascriptions," *Philosophical Studies* 39 (1981): 287–303, at pp. 294–295; R. M. Sainsbury, "On a Fregean Argument for the Distinctness of Sense and Reference," *Analysis* 43 (January 1983): 12–14; Takashi Yagisawa, Meaning and Belief, Ph.D. diss., Princeton University, 1981; J. Paul Reddam, Pragmatics and the Language of Belief, Ph.D. diss., University of Southern California, 1982. However, there are subtleties involved in Frege's Puzzle that these writers do not discuss. These subtleties will be developed in chapter 7 of this book with a new and stronger version of the puzzle, for which the solution presented here is simply irrelevant. (McKay comes very close to recognizing some of the finer aspects of the puzzle in his note 17, wherein he discusses an example (due to David Kaplan) involving a case of change of mind to suspension of judgment similar to the example to be presented in section 7.2 of this book. McKay's brief discussion of the example does not bring out the moral of my chapter 8.)
2. I am not talking here about overt verbal assent to a *sentence*, but about *mental* assent to a proposition.

The conception of belief as inward assent is apparently advanced by Saint Augustine in chapter 5 of *Predestination of the Saints*, where belief is analyzed as "to think with assent." For an illuminating contemporary discussion of the analysis of belief as assent to an entertained proposition, see H. H. Price, *Belief* (London: Allen and Unwin, 1969), especially series I, lectures 8 and 9, and series II, lectures 1–3.

In suggesting that belief might be understood in terms of inward assent, concurrence, or approval, I am not suggesting a reduction of belief to a phenomenological episode (in the style of Hume). By 'inward assent', etc., I do not mean merely a private, subjective experience directed toward or involving the relevant piece of information (such as the experience one typically has when reading or saying the words 'yes, I agree' to oneself, together with a "feeling of understanding" of these words, etc.). Such an analysis would be unacceptable for familiar philosophical reasons: In unusual circumstances someone could have these experiences without believing, and without even grasping, the proposition. By 'inward assent', etc., I mean a state of *cognition*, with everything that this entails.

Furthermore, in speaking of a *disposition* to inward assent, or other favorable *dispositions*, I do not mean merely an inclination, tendency, or propensity to assent, etc. (the usual philosophical use of 'disposition' as in, e.g., 'dispositional property'). Here again, it is possible for someone to have such inclinations without believing the proposition, and vice versa. In saying that someone is favorably "disposed" toward something, I mean that the person harbors a positive, favorable attitude (e.g. agreement, as opposed to disagreement or indifference) toward the thing. Typically, the harboring of this attitude will result in certain inclinations or propensities, but that is not part of the analysis of the attitude, and in extraordinary circumstances the harboring of the favorable attitude may not result in the typical inclinations. Conversely, the inclinations may be present in the absence of the favorable attitude. It is probably best to speak of 'attitudes' rather than 'dispositions'. So understood, the suggested analysis may appear unilluminating, but then at least it is not controversial.

I use the term 'believe', and its cognates, in such a way that one believes that *S* if one is convinced that *S*, of the opinion that *S*, confident that *S*, persuaded that *S*, etc., but it is not sufficient that one merely thinks it likely that *S*, guesses that *S*, suspects that *S* theorizes that *S*, assumes that *S*, or supposes that *S*.

3. I have in mind theories like that given by Robert Stalnaker in "Assertion," in *Syntax and Semantics 9: Pragmatics*, ed. P. Cole (New York: Academic, 1978). See also Stalnaker, "Indexical Belief," *Synthèse* 49, no. 1 (1981): 129–151; David Lewis, "What Puzzling Pierre Believes," *Australasian Journal of Philosophy* 59, no. 3 (1981): 283–289. The favored theory of propositions here is one that identifies propositions with sets of possible worlds—a theory on which propositions are even more coarse-grained than on the modified naive theory—though that is largely irrelevant to the main idea behind Stalnaker's account. Stalnaker claims that, in at least some propositional-attitude contexts, a 'that'-clause, ⌜that *S*⌝, will sometimes refer not to the proposition expressed by the sentence *S* but instead to a related proposition, which Stalnaker calls 'the diagonal proposition of the propositional concept for *S*'. This so-called diagonal proposition, if it is a proposition at all, is best identified as the singular proposition about *S* that it is true—or, more accurate, as the proposition *that the proposition semantically encoded by S, as uttered in a context, is true*. (Stalnaker shows reluctance to so identify the diagonal proposition. The coarse-grainedness of his favored theory of propositions enables him to avoid specifying the relevant proposition in this way; however, from the point of view of a more fine-grained theory (such as the modified naive theory or the Fregean theory), the metatheoretic proposition that *S* is true is the most plausible candidate for being the diagonal proposition.) I shall

not argue the point fully here. For present purposes, it is sufficient that this be one way of understanding what Stalnaker means by 'the diagonal proposition'. In effect, then, on Stalnaker's theory a 'that'-clause ⌜that S⌝ may be ambiguous. It sometimes refers to the proposition encoded by S, and it sometimes refers to a different, metatheoretic proposition about S itself. Rather than postulate this sort of complexity or ambiguity in connection with 'that'-clauses, it would be more plausible to claim that, in some cases, the speaker reporting a propositional attitude strictly speaking misspoke and, for complete accuracy, should have used a more complicated formal-mode 'that'-clause in place of the material-mode 'that'-clause used.

4. This consequence of the modified naive theory concerning nesting of propositional-attitude operators often goes unnoticed. In an attempt to soften the blow of the modified naive theory, it is sometimes argued that, for example, though the ancients strictly speaking did believe the proposition that Hesperus is Phosphorus, since this is just the trivial proposition that Hesperus is Hesperus, they did not realize that the proposition that Hesperus is Phosphorus is really the very same proposition as the trivial proposition that Hesperus is Hesperus, and hence they did not realize that they believed that Hesperus is Phosphorus. Similarly, it is sometimes argued that, since the name 'Hamlet' from Shakespeare's fiction actually refers to no one, there is no such thing as a proposition that Hamlet does not exist, and hence the sentence 'Hamlet does not exist' strictly speaking has no information content, but still there is a proposition that there exists no proposition that Hamlet does not exist, and it is true. All of this is inconsistent with the modifed naive theory. On the modified naive theory, if a is a single-word singular term (individual constant), then for any sentence ϕ_a containing a, barring quotation marks and other such aberrant devices, the 'that'-clause ⌜that ϕ_a⌝ refers to the singular proposition that is the information content of the sentence. It is tempting to think of the 'that'-term as a sort of description of the proposition by specifying its components, like ⌜the proposition made up of a and the property of being ϕ⌝, analogous to a set-theoretic abstraction term ⌜$(a$, the property of being $\phi)$⌝. But this is incorrect. A set-abstraction term ⌜$(\hat{x})\phi_x$⌝ may be regarded as a special sort of definite description, since it is equivalent to ⌜$(\iota y)[Set(y)$ & $(x)(x \in y \equiv \phi_x)]$⌝. Thus, a set-abstraction term is descriptional—specifically, descriptional in terms of the property of being a set with such-and-such membership. The 'that'-operator attaches to a sentence to form a singular term referring to the sentence's information content. Since 'that Plato is wise' refers to a different proposition from 'that the author of *The Republic* is wise', however, one cannot see the 'that'-term as referring to its referent proposition by mentioning the components of the referent proposition. Plato is not a component of the proposition that the author of *The Republic* is wise, though he is referred to by the component term 'the author of *The Republic*'. In a word, the 'that'-operator is nonextensional. One should think of the 'that'-operator as analogous to quotation marks, and of a 'that'-term ⌜that S⌝ as analogous to a quotation name, only referring to the information content of S rather than S itself. (See the introduction on the 'that'-operator.) A 'that'-clause, ⌜that ϕ_a⌝, then, is a singular term whose information value is the ordered pair of the information value of the 'that'-operator and the information content of ϕ_a, the letter being a singular proposition p about the referent of a. A sentence involving this 'that'-clause, ⌜ψ[that ϕ_a]⌝, encodes a singular proposition about the proposition p, to wit, that (the proposition identical with) it is ψ, and the 'that'-clause formed from this sentence, ⌜that ψ[that ϕ_a]⌝, refers to this singular proposition about p. If b is any proper name or other single-word singular term co-referential with a, then ⌜that ϕ_b⌝ refers to the very same proposition p, and ⌜ψ[that ϕ_b]⌝ encodes the same singular proposition about p that (the proposition identical with) it is ψ, so that ⌜that ψ[that ϕ_a]⌝ and ⌜that ψ[that ϕ_b]⌝ are co-

referential. In particular, if the sentence 'Jones realizes that he believes that Hesperus is Hesperus' is true, then what Jones realizes is a certain singular proposition about the proposition that Hesperus is Hesperus, to the effect that he believes it. Since the proposition that Hesperus is Hesperus is, according to themodified naive theory, the same proposition as the proposition that Hesperus is Phosphorus, another way of specifying what Jones realizes, according to the modified naive theory, is 'that Jones believes that Hesperus is Phosphorus'. Hence, on the modified naive theory, if the original sentence is true, so is 'Jones realizes that he believes that Hesperus is Phosphorus'. The proposition that Jones believes that Hesperus is Hesperus is the same proposition as the proposition that Jones believes that Hesperus is Phosphorus, and thus if Jones realizes the former he realizes the latter. Similarly, if the nonexistence of Hamlet means that there is no such proposition as the proposition that Hamlet does not exist, then it also means that there is no such proposition as the proposition that the proposition that Hamlet does not exist does not itself exist. Few philosophers—even direct-reference theorists who accede to the modified-naive-theoretical claim that the ancients strictly speaking believed that Hesperus is Phosphorus—have been willing to endorse these further consequences of the modified naive theory. Properly seen, however, they are no more unacceptble than the better-known controversial consequences of the modified naive theory.

These points concerning nested occurrences of 'that'-clauses are important in connection with the modified naive theory's account of Mates's problem concerning nested propositional attitude contexts. See appendix B.

Chapter 7

1. One special feature of Elmer's Befuddlement is that Elmer knows the relevant individual, Bugsy, by name, and by the same name in both of his guises, and Elmer comes to have his beliefs and his lack of belief concerning Bugsy (at least partly) by means of that name. This removes the wrongheaded temptation to identify the information value of a name with the name itself, for in this example Elmer's conflicting attitudes are directed toward a single sentence, 'Bugsy Wabbit is dangerous', involving a single name, 'Bugsy Wabbit'. Quine gives a name to his corresponding character, Ortcutt, but he frames the problem primarily in terms of definite descriptions—'the man in the brown hat' and 'the man seen at the beach'. This introduces a host of further issues, some of which have tended to distract commentators from the primary philosophical issues involved in situations involving ignorance of an identity. Kripke framed his original example concerning Pierre in London using names instead of descriptions, but he used two different syntactic shapes, 'London' and 'Londres', which correspond in the example to two different guises of the city, and which are correct translations of one another. This also has the unfortunate tendency to digress the course of the discussion toward a host of issues concerning translation—issues which are, and which Kripke recognizes to be, entirely irrelevant to the primary philosophical problems raised by the example. In this respect, Elmer's Befuddlement is more like Kripke's more pointed example concerning Paderewski (discussed here in note 4 to chapter 4).

   The most important aspect of Elmer's Befuddlement is the fact that Elmer has *changed his mind* about something and withholds belief where he once had an opinion. This aspect of the example—the change of mind from having an opinion to suspension of judgment—poses the most pressing and difficult philosophical problems, and is at the same time the most philosophically illuminating feature of the example. The importance of suspension of judgment to issues concerning propositional attitudes,

especially de re propositional attitudes, was first noticed by Kaplan in an important and underappreciated argument in section XI of "Quantifying In" [at pp. 139–142 of *Reference and Modality*, ed. L. Linsky (New York: Oxford University Press, 1971)]. See also Kripke, "A Puzzle About Belief," at p. 258. The similarity of my example involving Elmer and Bugsy, and the use to which I put it, to Kaplan's continuation of Quine's example should be apparent, though I do not argue for exactly the same conclusions as Kaplan. In particular, I do not argue, and I do not believe, that a de re or relational belief is reducible to a de dicto or notional one. (See note 7 to chapter 3.) Moreover, my proposal below to analyze belief as the existential generalization of a ternary relation among believers, propositions, and something else, unlike Kaplan's analysis in "Quantifying In" of de re belief, allows the formulation of de re belief in the manner of '$(\exists x)$[x is Elmer's friend & Elmer believes that x is dangerous]', with ordinary unrestricted, objectual quantification into the same 'believes' predicate used in the formulation of de dicto belief.

2. This contrasts with Kripke's puzzle about belief, where, on the modified naive theory, the subject comes to have incompatible beliefs. The naive theory's account of Kripke's puzzle is discussed below in appendix A.
3. Tyler Burge, in "Kaplan, Quine, and Suspended Belief," *Philosophical Studies* 31 (1977); 197–203, responds to Kaplan's argument concerning suspension of judgment (cited in note 1 above) by observing that suspension of judgment in the face of continued belief, such as Elmer's, can be expressed coherently by means of a certain conjunction—in our case, the conjunction of

   Elmer believes that Bugsy is dangerous

   with

   $(\exists \alpha)$[Elmer believes that Bugsy = α, but Elmer does not believe that α is dangerous],

   where 'α' ranges over Fregean individual concepts and Russellian intensional entities that are like Fregean individual concepts except for having nonintensional entities as constituents. (I have rephrased Burge's proposal extensively to frame it in the terminology of the present essay and to fit it to the example under discussion, under the presupposition of the correctness of the modified naive theory. I believe that I do not seriously misrepresent the spirit of Burge's original proposal, at least as far as the present discussion is concerned.) Burge goes on to argue that the second conjunct, expressing Elmer's withholding belief regarding his friend's dangerousness, is richer than it needs to be, and he suggests that Elmer's doxastic state can be correctly described by some simpler conjunction like 'Elmer believes Bugsy to be *this* man and that *this* man is dangerous, and Elmer also believes Bugsy to be *that* man, but Elmer does not believe that *that* man is dangerous', using each occurrence of a demonstrative with reference to Bugsy. This I regard as entirely unsatisfactory as a solution to the paradox, but I shall not argue the case here. In any event, this latter proposal clearly involves a rejection of the modified naive theory, and is therefore irrelevant to the present discussion.

   Burge's former proposal, as stated, also fails to solve the puzzle. If Elmer is sufficiently clever but unreflective, the second conjunct will be satisfied even if Elmer does not withdraw his former opinion about the criminal's being dangerous. For example, let α be the concept *the x such that x is Bugsy's mother's son's father's son, if Bugsy has no brothers, and x is Bugsy's father's son but not one of Bugsy's brothers, otherwise.* Elmer may know on April 1 that the person so characterized is Bugsy; however, he may never have entertained, and may have no favorable disposition toward, any other proposition involving this concept. This problem can be overcome by strengthening

Burge's proposed second conjunct, but even so, the general idea of expressing Elmer's suspension of judgment by way of Elmer's failure to believe some *other* related proposition gives insufficient recognition to the fact that, whatever other propositions Elmer may fail to believe, his recent change of heart and his present cautious attitude toward the very singular proposition about Bugsy that he is dangerous, when he takes it as a proposition about the notorious jewel thief, have all of the signs and trappings of withholding belief and suspension of judgment. His present attitude clearly would *constitute* withheld belief concerning that very proposition had Elmer not formed his earlier opinion about the dangerousness of his friend. Somehow this feature of Elmer's doxastic state—the fact that Elmer has withdrawn his earlier favorable disposition and now adopts a cautious "wait-and-see" attitude that, at least ordinarily, constitutes a suspension of belief toward the relevant singular proposition—must be expressed if we are to capture the gist of Elmer's complex state.

Chapter 8

1. Withheld belief, as defined here, is compatible with (in fact, perhaps entailed by) disbelief (belief of the negation). One can similarly define suspension of judgment so that its analysis is

   $(\exists x)[A \text{ grasps } p \text{ by means of } x \;\&\; \sim BEL(A, p, x) \;\&\; \sim BEL(A, \sim p, \text{Neg}(x))]$,

   e.g., under at least one relevant way of taking p, A is disposed neither to inward agreement nor to inward disagreement with respect to p. So understood, suspension of judgment entails withheld belief with respect to both the proposition in question and its negation, but not vice versa. The main idea is to see the various doxastic states of belief, disbelief, withheld belief, and suspension of judgment as existential generalizations of ternary relations relativized to guises, or some such items, so that it is consistent and reasonable for someone to be in conflicting doxastic states (e.g. belief and disbelief, or belief and suspension of judgment) with respect to the very same proposition.
2. There are important limitations to this device inherent in the complexities of natural language. Suppose that Jones believes that he is the best logician in the department, so that something like the following obtains:

   BEL[Jones, that he is the best logician in the department, *f*(Jones, 'I am the best logician in the department')],

   where the function $f(x,S)$ is something like the way x takes the proposition encoded by S *with respect to a context in which x is agent* when it is presented to him through the very sentence S. That is, Jones assents to the proposition that he is the best logician in the department when he takes it in the way he does when he presents it to himself through the sentence 'I am the best logician in the department'. It will not do in this case to use for the 'that'-clause the sentence that determines via the function f the way Jones takes the relevant proposition when agreeing to it, since Jones does not believe that *I* am the best logician in the department. It is quite possible that belief attributions of the form $\ulcorner a$ believes that he or she is $\phi\urcorner$, understood on the reflexive reading of the pronoun, involve the analogous cancelable implicature, suggestion, or presumption, expressible using '*BEL*' by $\ulcorner BEL[a$, that he or she is ϕ, $f(a$, 'I am ϕ')$]\urcorner$, and similarly for other so-called first-person propositional-attitude attributions (i.e., propositional-attitude attributions concerning oneself) and for other indexical or tensed propositional-attitude attributions. Thus, for example, if Jones believed at time t that the meeting was over by *then*, he probably did so by agreeing to that

information when taking it the way he would had it been presented to him at *t* through the present tensed sentence 'The meeting is over by now'. Where the relevant implicature, suggestion, or presumption is false, competent speakers may be inclined, erroneously from the point of view of truth, to deny the attribution, just as in the case of 'The ancient astronomer believes that Hesperus is Phosphorus' or 'The ancient astronomer believes that Hesperus appears at dawn'. This fact may help explain the widespread intuition, tapped by Hector-Neri Castañeda in support of his theory of so-called quasi-indicators and by others in support of equally or even more dramatically philosophical theses, that such attributions are literally false in such cases. See for example Castañeda, "Indicators and Quasi-Indicators," *American Philosophical Quarterly* 4, no. 2 (1967): 85–100; "On the Logic of Attributions of Self-Knowledge to Others," *Journal of Philosophy* 65, no. 15 (1968): 439–456; Roderick Chisholm, *The First Person* (Minneapolis: University of Minnesota Press, 1981); David Lewis, "Attitudes *De Dicto* and *De Re*," *Philosophical Review* 87 (1979): 513–443. For an account of first-person propositional-attitude attributions similar in broad outline and spirit to the one proposed here (although, as David Austin has pointed out, apparently lacking the full resources of the ternary account in terms of the *BEL* relation), see Steven Boër and William Lycan, "Who Me?," *Philosophical Review* 89, no. 3 (1980): 427–466.

Chapter 9

1. On the other hand, twin-earth considerations of the type considered in subsection 5.1.2 would suggest that the objects that serve as the third relatum for the *BEL* relation must be more fine-grained even than purely conceptual representations (Fregean senses) or the purely conceptual states of someone grasping such representations. For each time that a thinker is presented with a single proposition but does not recognize it as the same proposition as one encountered under different circumstances, there must be a new such object *x* by which he or she is familiar with the proposition.

   It is not surprising (see note 1 to chapter 7) that one direct-reference theorist who recognized the need to analyze belief as the existential generalization of a three-place relation among believers, propositions, and something else is David Kaplan. In "Demonstratives" (draft no. 2, section III and note 19.2) Kaplan proposes invoking characters as the third relatum. This suggestion, however, does not solve the problems posed by Elmer's Befuddlement, as Kaplan admits, since in Elmer's case the relevant characters corresponding to the two ways in which Elmer may believe that Bugsy is dangerous are one and the same.

   Kaplan's account inspired a similar account by John Perry. In "The Problem of the Essential Indexical" [*Noûs* 13 (1979): 3–21] Perry invokes so-called self-locating beliefs—indexically expressed beliefs concerning who one is, where one is, or when it is—to urge the acknowledgment of what he calls 'belief states' not individuated by the propositions they determine in a context. Perry's thesis appears to be that one believes a proposition *p* by being in a belief state that, in the believer's context, determines *p*, and that one can believe *p* by being in one belief state that determines *p* and yet fail to be in some other relevant belief state that also determines *p*. (Perry phrases his argument somewhat misleadingly by claiming, on pp. 11–12, that the phenomenon of self-locating beliefs "poses a problem for" an account of belief as a binary relation between believers and propositions, sometimes singular propositions. In fact, Perry's solution apparently preserves this binary relational account, and couples it with an existential analysis of belief in terms of belief states, emphasizing that the latter are crucial in psychological explanation.) Cases—such as Elmer's—of withholding

belief together with continued belief demonstrate that the general problems posed by self-locating beliefs and other beliefs formulated by means of an indexical are not peculiar to these special subcases of de re beliefs, and arise even with de re beliefs whose psychological explanatory force does not involve an "essential indexical." In fact, from a sufficiently broad point of view, the problems posed by indexical beliefs are essentially the same challenging questions that arise in Frege's Puzzle.

Perry's account can be applied to Elmer's Befuddlement: In April, Elmer believed that Bugsy is dangerous twice over, by being in both of two relevant belief states that determine this same proposition; by June 1, Elmer falls out of one of these belief states, and believes the proposition in only one way rather than two. This is, in effect, a restatement of the situation that gives rise to new challenging questions. Perry's account requires that belief states are not individuated by the propositions they determine, since the same belief state determines different propositions in different contexts and the same proposition is determined in different contexts by different belief states. Hence, a belief state is not to be confused with the state of having a certain belief. What, then, is a belief state, and how are belief states individuated? (Perry himself provides little in the way of a positive characterization of belief states, except that they are crucial for the explanation of behavior, and he allows on p. 20 that this is not even a sketchy account of belief.)

Following Kaplan's account, Perry claims (pp. 18–19) that a belief state might be characterized by a sentence which the believer "accepts" and which expresses the determined proposition in the believer's context, and that, where the characterizing sentence contains no indexicals, there is a "perfect correspondence" between the belief state and the determined proposition. Elmer's situation refutes this attempt to characterize belief states, since the two relevant belief states correspond to the same sentence, 'Bugsy Wabbit is always dangerous throughout his lifetime', which does not contain any indexicals or ambiguities. Belief states apparently must be more fine-grained than sentence meanings, and more so even than unambiguous sentences of the public language, if they are to deal adequately with Elmer's Befuddlement.

2. Thus, I disagree sharply with Stephen Schiffer's remarks in the opening passages of "The Basis of Reference" [*Erkenntnis* 13 (1978): 171–206], where he "locates the central question of the theory of reference in the theory of propositional attitude psychology." He adds that "the basis of a theory of reference must . . . be a theory of thought in the mind of a person using a singular term," and expresses dissatisfaction with "the cavalier way in which the 'Fido'-Fido theory ignores the connection between semantics and psychology" (pp. 171–175). On the conception I am endorsing here, the theory of reference and the theory of meaning generally are concerned with propositional-attitude psychology only to the extent that they are concerned with the *cognitive information content* of thought, and the modified naive theory provides an accurate and adequate picture of the structure of *semantically encoded* information content. The means by which a thinker is familiar with a thought content is perhaps a matter of pragmatics, or psychology proper, but it is not the business of the theory of reference, the theory of meaning, or semantics generally.

More recently, Schiffer has seemed more sympathetic to something like the original naive theory. See his "Indexicals and the Theory of Reference," *Synthèse* 49 (1981): 43–100. However, he apparently maintains that singular propositions about external objects cannot be the cognitive contents of beliefs, thoughts, and other attitudes, and he maintains that semantics is reducible to, and thus contained within, psychology. I disagree on both counts, but my concern here is to argue against only the former claim.

3. See note 7 to chapter 3. This is not to say that there are no philosophical problems involving de re modality or de re propositional attitudes. It may be very difficult to

decide whether someone has a certain de re belief in certain circumstances, or to state nontrivial conditions (e.g., involving certain de dicto beliefs) under which one has a certain de re belief. The modified naive theory does not make these questions disappear. They may be questions for the philosophy of psychology as much as for the philosophy of language. Analogously, the modified naive theory does not put to rest the notoriously difficult metaphysical problem of essentialism: What properties of a thing are essential to that thing? (See my *Reference and Essence*, part II.) These philosophical questions which remain even on the modified naive theory do not concern the intelligibility of quantification into modal or propositional-attitude contexts; they concern whether the truth conditions for these de re locutions are fulfilled.

4. See Kripke, "A Puzzle About Belief," pp. 245, 251–252.

Appendix A

1. Kripke, "A Puzzle About Belief," pp. 248–249. See p. 249 for Kripke's explanation of the force of the expressions 'normal English speaker', 'sincere', and 'reflection' as they occur in the disquotation principle.
2. Ibid., p. 249.
3. See Ruth Barcan Marcus, "A Proposed Solution to a Puzzle About Belief," in *Midwest Studies in Philosophy VI: The Foundations of Analytic Philosophy*, ed. French et al. (Minneapolis: University of Minnesota Press, 1981); "Rationality and Believing the Impossible," *Journal of Philosophy* 80, no. 6 (1983): 321–338. Although Marcus rejects Kripke's disquotation principle, she apparently would allow that sincere utterance of a sentence by a normal speaker does carry over into assertion (or "claim") by the speaker of the content of the sentence in his or her context. Marcus's rejection of the original disquotation principle is coupled with the surprising claim that it is impossible to believe what is impossible. Even one who shares Marcus's unusual view must acknowledge that it is easy to assert something impossible.
4. Kripke, "A Puzzle About Belief," pp. 266 and 269.
5. Questions concerning logic and logical attributes are generally thought to be clearer and somehow more concrete than questions concerning belief and the other propositional attitudes. Quine is relatively content with such notions as *logical truth* but harbors skepticism with respect to such notions as *relational belief*, which he sees as vague and context-relative. See W. V. O. Quine, "Intensions Revisited," in *Contemporary Perspectives in the Philosophy of Language*, ed. French et al. (Minneapolis: University of Minnesota Press, 1979). In fact, the logical notions involve a related can of worms. It is tempting to suppose that the logical notions are on firmer ground because they apply to sentences (in a language) rather than propositions. Consider then the English sentence 'Aristotle was Aristotle'. Is it a logical truth (a valid theorem)? Well, strictly speaking it is ambiguous. Let us disambiguate it. Suppose that the first occurrence of 'Aristotle' refers to the philosopher and the second to the shipping magnate. Then it is surely not a logical truth, since it is false. Of course, if instead both occurrences of 'Aristotle' had referred to the philosopher, then the expression form would have been a logical truth. Now, suppose once again that, though we do not know it, the philosopher of antiquity did not die in ancient times as we think, but went into hiding, discovered the philosopher's stone which slows down the aging process, and emerged in the twentieth century as the powerful shipping magnate. What then of the logical status of the disambiguated sentence in question? Can it be a logical truth, with the first occurrence of 'Aristotle' intended to apply to the philosopher and the second occurrence to the shipping magnate, even though no one now alive who understands the sentence as it is intended is in a position even in principle to recognize

it as a logical truth? Can a sentence be a simple logical truth of the form $\ulcorner a = a \urcorner$ even though an ideally perfect logician who fully understands the sentence would have to reckon it not logically valid? (An ideally perfect logician would, of course, believe with the rest of us that Aristotle the philosopher is not the same person as Aristotle the shipping magnate. Perfect logical acumen is not the same thing as omniscience.) Or is it not really logically valid after all? If so, what differentiates it from the disambiguated sentence with both occurrences of 'Aristotle' intended to apply to the same person? Surely the latter sentence is a logical truth. How can the former, if true, differ at all in logical status from the latter? Thus, some version of Frege's Puzzle arises not only with our concept of information, but also in a particularly sharpened form with the various logical notions of logical truth, entailment, consistency, and so on. (One of the things that the example suggests is that such notions as that of logical truth should not be defined or analyzed in terms of what an ideal logician would reckon concerning the truth value or theoretical status of a sentence. See also section 5.4 above.)

6. Kripke, "A Puzzle About Belief," p. 257.
7. David Lewis, "What Puzzling Pierre Does Not Believe," *Australasian Journal of Philosophy* 59, no. 3 (1981): 283–289, at pp. 288–289. See also Marcus, "A Proposed Solution to a Puzzle About Belief," pp. 504–508.

   Lewis is not explicitly concerned with the modified naive theory and its solution to Kripke's puzzle, but more generally with a "conception of belief and its objects" according to which "when we characterize the content of belief by assigning propositional (or other) objects, . . . [we are] characterizing partly the believer's inner state, partly the relations of that state to the outer world" (p. 288). It is on this conception of belief, according to Lewis, that Kripke's puzzle vanishes. It is a familiar fact, at least since Putnam's twin-earth thought experiment, that any minimally plausible theory (including the modified naive theory) on which the objects of belief are true or false (or neither) irrespective of context employs this conception of belief as generally not merely an inner, psychological state (in the narrow, "pure" sense). Two different individuals in different contexts but in the same (purely) psychological state will generally have different beliefs; they will generally be in differing doxastic states. In attributing to Elmer a belief that Bugsy Wabbit is dangerous, we cannot be merely characterizing Elmer's inner, psychological state, since a counterpart of Elmer on some other planet can be in the very same inner state as Elmer with respect to a counterpart of Bugsy on the other planet, though Bugsy is indeed dangerous whereas his alien counterpart is innocuous. We cannot describe Elmer's counterpart as also believing that Bugsy Wabbit is dangerous, since it is true that Bugsy Wabbit is dangerous but what Elmer's counterpart believes is false. (See also Wittgenstein, *Philosophical Investigations* II, xi, p. 217.)

   Lewis erroneously claims in this connection that "as soon as we accept the [compatibility of the subject's] beliefs as a datum—[as Kripke invites us to do]—we are committed to the narrowly psychological conception of belief and its objects" (p. 289). The compatibility of the subject's beliefs can be maintained together with a conception of belief as not generally a wholly internal state by denying Elmer one or both of his incompatible beliefs. Lewis's preferred narrowly psychological conception, in effect, denies Elmer both of his beliefs. By contrast, the theory I am advocating allows Elmer both of his beliefs without imputing to him any logical blunder. The Putnamian considerations mentioned here effectively preclude Lewis's preferred conception of belief as an option.
8. Kripke, "A Puzzle About Belief," p. 258.
9. Though not exactly for the reasons urged by Marcus ("A Proposed Solution," p. 509). Marcus misunderstands the quotation principle as asserting that, for every

proposition a believer believes, there is some sentence which encodes the proposition and to which the believer would assent. Marcus points out that this is falsified by infants and higher animals. But she is wrong to think that this, in itself, refutes Kripke's quotation principle, which is restricted to (normal) English speakers. The "principle" is actually a principle *schema*, and hence it is actually infinitely many principles that are automatically restricted to beliefs whose content is encoded by some English sentence, a sentence which a (normal) English speaker may be expected to understand fully.

Appendix B

1. Kaplan has argued convincingly that certain indexical sentences, e.g. 'I exist' and 'I am here now', are logically valid, since, no matter what the range of possible contexts, they are true with respect to every possible context. This result yields the consequence that such propositions as that Kaplan exists in A.D. 2100 and that Kaplan is in Budapest in A.D. 2100 are logically valid with respect to certain possible contexts (in which they are true). Even so, these propositions are not logically valid simpliciter, as the derivative concept is defined here, since they are not logically valid with respect to every possible context and time. It seems entirely appropriate that such propositions as these are not counted logically valid simpliciter, even if there are logically valid sentences capable of encoding them, since these propositions need not even be true in the actual world.

   There is an alternative and more direct way to define logical validity for propositions on the modified naive theory. Since propositions are similar in structure to the sentences that encode them, it is possible to construct a model-theoretic analysis of propositions, treating individual constituents as if they were atomic singular terms, n-ary relations as if they were n-place predicates, and so on. We define the notion of a *p-structure* by analogy to a model structure for a modal and temporal (but nonindexical) language. A proposition is *logically valid*, in this model-theoretic sense, if it is true in every p-structure. This alternative conception of logical validity is developed in appendix C below.
2. In light of the preceding note, this may include indexical logically valid sentences, e.g. 'I exist' and 'I am here now', whose contents are not themselves logically valid. Indeed, sentences like these (especially 'I exist') have been put to some of the same philosophical uses, e.g. in epistemological foundationalism, as the classical logical validities.
3. Quine, "Two Dogmas of Empiricism," in his *From a Logical Point of View* (New York: Harper and Row, 1953), pp. 22–23. I would urge consideration of the modified naive theory as answering some (though not all) of Quine's well-known worries concerning the concept of analyticity and the concept of information content on which analyticity depends.
4. Traditional examples of analyticity include 'Bachelors are unmarried', 'Vixen are foxes', and 'Brothers are siblings'. It is tempting to suppose that all such examples (excluding genuinely logically valid sentences) are deemed synthetic on the modified naive theory. Whether such sentences are analytic depends on whether 'bachelor' is synonymous with some phrase including the term 'unmarried' (e.g. 'unmarried man eligible for marriage'), whether 'vixen' is synonymous with 'female fox', and so forth. On the modified naive theory, the infomation value of a simple expression such as 'bachelor' or 'vixen' is unlike that of a compound expression whose information value is a complex made up of the information values of the compound's information-valued components. Thus, it would seem that 'bachelor' is not strictly synonymous

with any compound expression involving 'unmarried' as a component, that 'vixen' is not strictly synonymous with any compound expression involving 'fox' as a component, and so on. (This idea will be developed further in this appendix.) Common noun phrases, however, may not be typical compound expressions. The predicate phrase 'is a female fox' is at least logically equivalent to, and perhaps even synonymous with, the semi-formal '$(\lambda x)(x$ is female & x is a fox) (or 'x, x is female and x is a fox'), whose information value with respect to any time t is plausibly identified with the property of being a female fox at t. (See note 4 to chapter 1.) This property is equally a plausible candidate for being the information value of 'vixen' with respect to t. Thus, 'Vixen are foxes' would indeed have the same meaning as (something logically equivalent to) the logically valid 'Female foxes are foxes'; similarly for 'Bachelors are unmarried' and so on.

5. Kripke, *Naming and Necessity*, p. 101.
6. Ibid., pp. 102–103.
7. Ibid., p. 104.
8. Ibid., pp. 108–109. See also the passages quoted in section 6.1 above.
9. I am not arguing here that there are no (nontrivial) examples of necessary a posteriori sentences, only that, contra Kripke, sentences like 'Hesperus, if it exists, is Phosphorus' are not of this sort, because they are not a posteriori but a priori (in fact, analytic). Other, more interesting examples alleged by Kripke to be necessary yet a posteriori are 'Water, if there is any, is H_2O', 'Gold, if there is any, has atomic number 79', 'Tigers, if there are any, are mammals', and 'Elizabeth, if she exists, sprang from the particular gametes s and e' (assuming the last is true). These are indeed a posteriori. I also agree that they—or at least certain weakened replacements, e.g. 'Water is not CO_2'—are necessary. For a detailed analysis of the relation between the question of the modal status of these sentences and Kripke's direct-reference theory, see my *Reference and Essence*, part II.
10. Alonzo Church, "Intensional Isomorphism and Identity of Belief," *Philosophical Studies* 5, no. 5 (1954): 65–73, at p. 68. Church allows in a footnote that the expressions 'fortnight' and 'period of fourteen days' may not be exactly synonymous in English, suggesting instead that 'fortnight may be equivocal in meaning between 'period of fourteen days' and 'period of two weeks'. It is clear, however, that Church believes both that there is no reason in principle why a single word of English should not in general be exactly synonymous with a compound expression of English and that the meaning in English of 'fortnight' in particular is at least very much like that of 'period of fourteen days'.
11. One usually neglected complication here is that the term 'meter', in contrast with other examples of what is alleged to be the same phenomenon, is a common noun rather than a proper name and hence apparently a general term rather than a singular term. The application of the notion of reference (or "designation") to general terms is unclear at best and perhaps even an error or a confusion. However, the grammatical category of the term 'meter' need not be taken as definitively determining the term's semantic or logical category, for, in contrast with the case with other common nouns, numerical phrases formed by means of the word 'meter' are sometimes used as singular terms (referring to lengths or linear distances, as in 'The length of S is three meters') and apparently never used as monadic quantifiers ('Three meters are . . .'). For present purposes, I propose to treat unit measurement terms such as 'meter' as singular terms that combine with number terms in combination with pluralization ('three meters') to form new singular terms.
12. Kripke, *Naming and Necessity*, pp. 54–57.
13. That they do not might be argued via propositional attitudes: A particular prisoner who has been confined for exactly fourteen days with no clocks or windows that

might indicate the passage of time estimates concerning *the length of his period of confinement so far* that it is at least fifteen days. He may know fully well that the word 'fortnight' refers to the length of time fourteen days, or he may have no idea what the word 'fortnight' means. He need not express his estimate concerning the length of his period of confinement using the word 'fortnight'. Nevertheless, he has in fact formed a guess concerning the length of time one fortnight, and we may correctly (albeit misleadingly) report him as guessing that a fortnight is at least fifteen days. (Be very careful: it would be incorrect, of course, to report him as having formed any belief or guess concerning the *word* 'fortnight'. Our correct report must not be so misconstrued. His guess concerns a fortnight, not 'a fortnight'.) Still, he does not guess that fourteen days is at least fifteen days. (To be sure, this argument is not overpowering. This is because it is difficult to resist the temptation to read the locution ⌜*a* believes that . . . a fortnight . . .⌝ as ascribing a belief concerning the word 'fortnight' rather than one concerning the length of time, a fortnight.)

Considering Church's original argument regarding Mates's problem, it is indeed possible for someone to believe of the length of time fourteen days that the seventh consulate of Marius lasted less than that, and hence to believe that the seventh consulate of Marius lasted less than a fortnight, without believing that the seventh consulate of Marius lasted less than a period of fourteen days. Despite my difference with Church concerning the particular example of 'fortnight' and 'period of fourteen days', I am in general agreement with his analysis of Mates's problem. From the point of view of the modified naive theory (which Church rejects), a better example for Church's purpose is provided by 'Cicero'-'Tully' cases. Mates, "whatever he himself may tell us," believes that whoever believes that Cicero is an orator believes that Tully is an orator—even if it is also true that he doubts that (withholds belief concerning whether) whoever believes that Cicero is an orator believes that Tully is an orator. Despite his doubts, he also believes that whoever believes that Cicero is an orator believes that Cicero is an orator, and that is all it takes. Compare Elmer's Befuddlement.

14. D. Kaplan, "Bob and Carol and Ted and Alice," in *Approaches to Natural Language*, ed. Hintikka et al. (Dordrecht: Reidel, 1973), at p. 501).
15. By and large, Fregeans, like Church, suppose the former. Keith Donnellan, a pioneering direct-reference theorist, supposes the latter in "The Contingent A Priori and Rigid Designators," in *Contemporary Perspectives in the Philosophy of Language*, ed. French et al. (Minneapolis: University of Minnesota Press, 1979), at p. 49.
16. The suggestion is apparently not one that Kaplan intended, or accepts.
17. Kripke, *Naming and Necessity*, p. 63. Kripke admits that the claim seems implausible, and that some version or variant of its denial may be true, but as far as I can tell he makes the claim all the same.
18. Donnellan argues forcefully along these lines in "The Contingent A Priori and Rigid Designators," especially on pp. 52–58. His arguments are based, to some extent, on the contention that a reference-fixing definition by means of a definite description does not place the agent in a position to have any propositional attitudes, and hence to have knowledge (a priori or not), directly concerning the referent.
19. Here again, I am not arguing that there are no examples of contingent a priori sentences, only that, contra Kripke, the particular sentence 'Stick *S*, if it exists, is exactly one meter long at *t*' is not of this sort because it is not a priori but a posteriori, even for someone who introduces the term 'meter' as a term for the length of stick *S* at *t*. Similarly, to use an example of Kaplan's, even if one were to introduce the term 'Newman 1' as a name for the first child to be born in the twenty-second century, the sentence 'If anyone is born in the twenty-second century before anyone else, then it will be Newman 1' is not contingent a priori, because it is not a priori,

although it may be a priori for the introducer that this sentence is true in his or her idiolect. Or again, most of us are in no position to know or even to believe anything at all de re concerning whoever happens to be the world's shortest spy, let alone anything a priori. Introducing a name for him does nothing to change this, even if doing so does place one in a position to be able to *assert* something de re concerning the shortest spy using his newly introduced name. See Kaplan, "Quantifying In," at pp. 126–127 and 135 of Linsky 1971. (But see also Kaplan, "Dthat," at p. 397 of French et al. 1979.)

I agree with Kripke, however, that there are contingent a priori sentences. If *S* is any contingent false sentence, then the logically valid conditional 'If *S*, then actually *S*' is a trivial example of this sort. The sentence 'I am here now' may be another. (See note 1.) More important, Kripke is also correct that fixing the referent of a term by means of a description in some cases yields certain nontrivial examples of contingent a priori sentences—roughly, whenever the description is both nonrigid and such that merely grasping its information value ipso facto places the user in a position to form de re beliefs concerning the referent qua the thing so described. This situation obtains in certain cases where the description invokes certain sorts of causal contact or certain other contextual relations involving the referent. (See note 3 to chapter 9.) Kripke has argued, in conversation, that even his 'meter' example can be bolstered or restored with a suitably guarded description, e.g. 'the length of the stick visually veridically presented to me by this very visual perception', spoken to oneself while one is introspectively ostending a particular visual perception (assuming that the perception in question might have visually veridically presented a stick of imperceptibly different length, so that the description is nonrigid). It is not clear, however, that the 'meter' example can be thus restored. One can know a priori concerning a particular stick *S* that if such-and-such a visual perception is visually veridically presenting a certain stick to oneself, then *S* is that stick. Perhaps this bit of a priori knowledge is even contingent. But is it possible to know concerning a particular length—a priori and without measurement—that if the visual perception in question is visually veridically presenting a stick to oneself, then the stick thus seen is exactly *that* long? If so, then it would seem that we can know just by looking, without bothering to measure, the *exact* size of any seen (middle-size) physical object provided we are not visually deceived.

Bibliography

Ackerman, D. "Proper Names, Propositional Attitudes and Non-Descriptive Connotations." *Philosophical Studies* 35, no. 1 (1979): 55–69.

Augustine. "To Believe Is to Think With Assent." Chapter 5 of *Predestination of the Saints.* In *Basic Writings of Saint Augustine,* volume 1, ed. W. J. Oates (New York: Random House, 1948).

Baker, L. R. "Underprivileged Access." *Noûs* 16 (May 1982): 227–241.

Barwise, J., and J. Perry. *Situations and Attitudes.* Cambridge, Mass.: MIT Press, 1983.

Biro, J. I. "What's in a Belief?" *Logique et Analyse* 27 (September 1984): 267–282.

Blackburn, S. "The Identity of Propositions." In *Meaning and Reference,* ed. S. Blackburn (Cambridge University Press, 1975).

Boër, S., and W. Lycan. "Who Me?" *Philosophical Review* 89, no. 3 (1980): 427–466.

Bradley, R., and N. Swartz. *Possible Worlds: An Introduction to Logic and Its Philosophy.* Indianapolis: Hackett, 1979.

Burge, T. "Kaplan, Quine, and Suspended Belief." *Philosophical Studies* 31 (1977): 197–203.

Burge, T. "Belief *De Re.*" *Journal of Philosophy* 69 (May 18, 1977): 338–362.

Burge, T. "Belief and Synonymy." *Journal of Philosophy* 75 (March 1978): 119–138.

Burge, T. "Self-Reference and Translation." In *Translation and Meaning,* ed. F. Guenthner and M. Guenthner-Reutte (London: Duckworth, 1978).

Burge, T. "Individualism and the Mental." In French et al. 1979.

Cappio, J. "Russell's Philosophical Development." *Synthèse* 46 (1981): 185–205.

Carnap, R. *Meaning and Necessity: A Study in Semantics and Modal Logic.* University of Chicago Press, 1947.

Carnap, R. "Reply to Leonard Linsky." *Philosophy of Science* 16, no. 4 (1949): 347–350.

Cartwright, R. "Propositions." In *Analytical Philosophy,* first series, ed. R. Butler (Oxford: Blackwell, 1966).

Castañeda, H. "Indicators and Quasi-Indicators." *American Philosophical Quarterly* 4, no. 2 (1967): 85–100.

Castañeda, H. "On the Logic of Attributions of Self-Knowledge to Others." *Journal of Philosophy* 65, no. 15 (1968): 439–456.

Chien, A. J. "Demonstratives and Belief States." *Philosophical Studies* 47 (1985): 271–289.

Chisholm, R. *The First Person: An Essay on Reference and Intentionality.* Minneapolis: University of Minnesota Press, 1981.

Church, A. "Review of Carnap's *Introduction to Semantics.*" *Philosophical Review* 52 (1943): 298–304.

Church, A. "On Carnap's Analysis of Statements of Assertion and Belief." *Analysis* 10, no. 5 (1950): 97–99. Also in Linsky 1971.

Church, A. "A Formulation of the Logic of Sense and Denotation." In *Structure, Method and Meaning: Essays in Honor of Henry M. Sheffer,* ed. Henle, Kallen, and Langer (New York: Liberal Arts Press, 1951).

Church, A. "Intensional Isomorphism and Identity of Belief." *Philosophical Studies* 5, no. 5 (1954): 65–73. Also in Salmon and Soames (forthcoming).

Church, A. *Introduction to Mathematical Logic I.* Princeton University Press, 1956.

Church, A. "Outline of a Revised Formulation of the Logic of Sense and Denotation." Part I: *Noûs 7* (March 1973): 24–33; part II: *Noûs* 8 (May 1973): 135–156.

Church, A. "A Remark Concerning Quine's Paradox About Modality." Spanish translation in *Analisis Filosofico* 2, nos. 1–2 (1982): 25–34; in English in Salmon and Soames (forthcoming).

Cresswell, M. J. *Structured Meanings: The Semantics of Propositional Attitudes.* Cambridge, Mass.: MIT Press, 1985.

Devitt, M. *Designation.* New York: Columbia University Press, 1981.

Donnellan, K. "Reference and Definite Descriptions." *Philosophical Review* 75 (July 1966): 647–658.

Donnellan, K. "Proper Names and Identifying Descriptions." In *Semantics of Natural Language,* ed. D. Davidson and G. Harman (Dordrecht: Reidel, 1972).

Donnellan, K. "Speaking of Nothing." *Philosophical Review* 83 (January 1974): 3–31. Also in Schwartz 1977.

Donnellan, K. "The Contingent *A Priori* and Rigid Designators." In French et al. 1979.

Donnellan, K. "Kripke and Putnam on Natural Kind Terms." In *Knowledge and Mind,* ed. C. Ginet and S. Shoemaker (Oxford University Press, 1983).

Dummett, M. *Frege: Philosophy of Language.* New York: Harper and Row, 1973.

Dummett, M. *Truth and Other Enigmas.* Cambridge, Mass.: Harvard University Press, 1978.

Eliot, T. S. "The Naming of Cats." In Eliot's *Old Possum's Book of Practical Cats.* New York: Harcourt, Brace, 1939.

Evans, G. "Pronouns, Quantifiers and Relative Clauses (I)." In *Reference, Truth and Reality,* ed. M. Plattes (London: Routledge and Kegan Paul, 1980).

Fine, K. "Properties, Propositions and Sets." *Journal of Philosophical Logic* 6 (June 1977): 135–192.

Frege, G. "Begriffsschrift." In English in Frege's *Conceptual Notation and Related Articles* (New York: Oxford University Press, 1972); also in Frege 1952.

Frege, G. "Function und Begriff." In English in Frege 1984; also in Frege 1952.

Frege, G. "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und Philosophische Kritik* 100 (1893): 25–50. In English in Frege 1984; also in Frege 1952.

Frege, G. "Über Begriff und Gegenstand." In English in Frege 1984; also in Frege 1952.

Frege, G. *Translations from the Philosophical Writings,* tr. P. Geach and M. Black. Oxford: Blackwell, 1952.

Frege, G. "Gedankengefüge." In English in Frege 1984; also in Frege 1977.

Frege, G. *Logical Investigations.* New Haven: Yale University Press, 1977.

Frege, G. "Logic." In Frege 1979.

Frege, G. *Posthumous Writings,* ed. H. Hermes, F. Kambartel, and F. Kaulbach, tr. P. Long and R. White. University of Chicago Press, 1979.

Frege, G. *Philosophical and Mathematical Correspondence,* ed. G. Gabriel, H. Hermes, F. Kambartel, C. Thiel, and A. Veraart, abr. B. McGuinness, tr. H. Kaal. University of Chicago Press, 1980.

Frege, G. *Collected Papers on Mathematics, Logic, and Philosophy,* ed. B. McGuinness, tr. M. Black, V. H. Dudman, P. Geach, H. Kaal, E.-H. W. Kluge, B. McGuinness, and R. H. Stoothoff. Oxford: Blackwell, 1984.

French, P., T. Uehling, and H. Wettstein, eds. *Contemporary Perspectives in the Philosophy of Language.* Minneapolis: University of Minnesota Press, 1979.

French, P., T. Uehling, and H. Wettstein, eds. *Midwest Studies in Philosophy V: Studies in Epistemology.* Minneapolis: University of Minnesota Press, 1980.

French, P., T. Uehling, and H. Wettstein, eds. *Midwest Studies in Philosophy VI: the Foundations of Analytic Philosophy*. Minneapolis: University of Minnesota Press, 1981.

Gödel, K. "Russell's Mathematical Logic." In *The Philosophy of Bertrand Russell*, ed. P. A. Schilpp (New York: Tudor, 1946).

Grice, H. P. "Logic and Conversation." In *The Logic of Grammar*, ed. D. Davidson and G. Harman (Encino: Dickenson, 1975).

Hacking, I. "Comment on Wiggins." In *Philosophy of Logic*, ed. S. Körner (Berkeley: University of California Press, 1976).

Harman, G. "How to Use Propositions." *American Philosophical Quarterly* 14, no. 2 (1977): 173–176.

Hintikka, J. *Knowledge and Belief*. Ithaca: Cornell University Press, 1962.

Hintikka, J. "Semantics for Propositional Attitudes." In *Philosophical Logic*, ed. J. W. Davis, D. J. Hockney, and W. K. Wilson (Dordrecht: Reidel, 1969). Also in Linsky 1971 and Hintikka 1969.

Hintikka, J. *Models for Modalities*. Dordrecht: Reidel, 1969.

Hintikka, J. *The Intensions of Intensionality and Other New Models for Modalities*. Dordrecht: Reidel, 1975.

Hodes, H. "The Composition of Fregean Thoughts." *Philosophical Studies* 41 (1982): 161–178.

Kamp, H. "Formal Properties of 'Now'." *Theoria* 37 (1972): 227–273.

Kaplan, D. "Quantifying In." In *Words and Objections: Essays on the Work of W. V. Quine*, ed. D. Davidson and G. Harman (Dordrecht: Reidel, 1969). Also in Linsky 1971.

Kaplan, D. "Dthat." In *Syntax and Semantics 9: Pragmatics*, ed. P. Cole (New York: Academic, 1978). Also in French et al. 1979.

Kaplan, D. "Bob and Carol and Ted and Alice." In *Approaches to Natural Language*, ed. J. Hintikka, J. Moravcsik, and P. Suppes. (Dordrecht: Reidel, 1973).

Kaplan, D. "On the Logic of Demonstratives." In French et al. 1979; also in Salmon and Soames (forthcoming).

Kaplan, D. "How to Russell a Frege-Church." *Journal of Philosophy* 72 (November 6, 1975): 716–729. Also in *The Possible and the Actual*, ed. M. Loux (Ithaca: Cornell University Press, 1979).

Kaplan, D. "Demonstratives" (draft number 2). Unpublished manuscript, Department of Philosophy, UCLA, 1977.

Kaplan, D. "Opacity." In *The Philosophy of W. V. O. Quine*, ed. P. A. Schilpp (forthcoming).

Kneale, W., and M. Kneale. "Propositions and Time." In *G. E. Moore: Essays in Retrospect*, ed. A. Ambrose and M. Lazerowitz (London: Allen and Unwin, 1970).

Kripke, S. "Identity and Necessity." In *Identity and Individuation*, ed. M. Munitz (New York University Press, 1971). Also in Schwartz 1977.

Kripke, S. *Naming and Necessity*. Cambridge, Mass.: Harvard University Press, 1980. Also in *Semantics of Natural Language*, ed. D. Davidson and G. Harman (Dordrecht: Reidel, 1972).

Kripke, S. "Speaker's Reference and Semantic Reference." In French et al. 1979.

Kripke, S. "A Puzzle About Belief." In *Meaning and Use*, ed. A. Margalit (Dordrecht: Reidel, 1979).

Kvart, I. "Kripke's Belief Puzzle." *Philosophy Research Archives* 9 (1983): 369–412.

Lewis, D. "Attitudes *De Dicto* and *De Re*." *Philosophical Review* 87 (1979): 513–543. Also in Lewis 1983, with postscripts.

Lewis, D. *Philosophical Papers I*. New York: Oxford University Press, 1983.

Lewis, D. "What Puzzling Pierre Believes." *Australasian Journal of Philosophy* 59, no. 3 (1981): 283–289.

Linsky, L., ed. *Reference and Modality*. New York: Oxford University Press, 1971.

Linsky, L. *Names and Descriptions.* University of Chicago Press, 1977.

Linsky, L. *Oblique Contexts.* University of Chicago Press, 1983.

Marcus, R. B. "The Identity of Individuals in a Strict Functional Calculus of Second Order." *Journal of Symbolic Logic* 12 (1947): 12–15.

Marcus, R. B. "Modalities and Intensional Languages." *Synthèse* 27 (1962): 303–322.

Marcus, R. B. "A Proposed Solution to a Puzzle About Belief." In French et al. 1981.

Marcus, R. B. "Rationality and Believing the Impossible." *Journal of Philosophy* 80, no. 6 (1983): 321–338.

Mates, B. "Synonymity." *University of California Publications in Philosophy* 25 (1950). Also in *Semantics and the Philosophy of Language*, ed. L. Linsky (University of Illinois Press, 1952).

Mill, J. S. *A System of Logic.* New York: Harper Brothers, 1843.

McGinn, C. "The Structure of Content." In *Thought and Object*, ed. A. Woodfield (New York: Oxford University Press, 1982).

McKay, T. "On Proper Names in Belief Ascriptions." *Philosophical Studies* 39 (1979): 287–303.

Montague, R. "The Proper Treatment of Quantification in Ordinary English." In Montague 1974.

Montague, R. *Formal Philosophy*, ed. R. Thomason. New Haven: Yale University Press, 1974.

Moore, G. E. *The Commonplace Book 1919–53*, ed. C. Lewy. New York: Macmillan, 1962.

Moore, G. E. "Facts and Propositions." *Aristotelian Society, Supplementary Vol. 7* (1927). Also in Moore's *Philosophical Papers* (New York: Collier Books, 1962).

Over, D. E. "On Kripke's Puzzle." *Mind* 92 (1983): 253–256.

Perry, J. "Frege on Demonstratives." *Philosophical Review* 86 (1977): 474–497.

Perry, J. "The Problem of the Essential Indexical." *Noûs* 13 (1979): 3–21.

Perry, J. "Belief and Acceptance." In French et al. 1980.

Perry, J. "A Problem About Continued Belief." *Pacific Philosophical Quarterly* 61 (1980): 317–332.

Plantinga, A. "The Boethian Compromise." *American Philosophical Quarterly* 15 (April 1978): 129–138.

Plato. *Sophist.* In *The Collected Dialogues of Plato*, ed. E. Hamilton and H. Cairns (Princeton University Press, 1961).

Price, H. H. *Belief* (London: Allen and Unwin, 1969.

Prior, A. N. *Past, Present and Future.* New York: Oxford University Press, 1967.

Prior, A. N. " 'Now'." *Noûs* 2 (1968): 101–119.

Prior, A. N. *Papers on Time and Tense.* New York: Oxford University Press, 1968.

Prior, A. N. *Objects of Thoughts.* New York: Oxford University Press, 1971.

Putnam, H. "Synonymity, and the Analysis of Belief Sentences." *Analysis* 14, no. 5 (1954): 114–122.

Putnam, H. "Meaning and Reference." *Journal of Philosophy* 70 (November 8, 1973): 699–711. Also in Schwartz 1977.

Putnam, H. "The Meaning of 'Meaning'." In *Minnesota Studies in the Philosophy of Science VII: Language, Mind, and Knowledge*, ed. K. Gunderson (Minneapolis: University of Minnesota Press, 1975). Also in Putnam's *Philosophical Papers II: Mind, Language, and Reality* (Cambridge University Press, 1975).

Putnam, H. "Comments [on Kripke 1979]." In *Meaning and Use*, ed. A. Margalit (Dordrecht: Reidel, 1979).

Quine, W. V. O. *Mathematical Logic.* New York: Norton, 1940.

Quine, W. V. O. "Two Dogmas of Empiricism." *Philosophical Review*, January 1951. Also in Quine 1953.

Quine, W. V. O. "Reference and Modality." In Quine 1953; also in Linsky 1971.
Quine, W. V. O. *From a Logical Point of View*. New York: Harper and Row, 1953.
Quine, W. V. O. "Quantifiers and Propositional Attitudes." *Journal of Philosophy* 53, no. 5 (1956): 177–187. Also in Quine's *The Ways of Paradox* (New York: Random House, 1966) and in Linsky 1971.
Quine, W. V. O. *Word and Object*. Cambridge, Mass.: MIT Press, 1960.
Quine, W. V. O. "Intensions Revisited." In French et al. 1979.
Reddam, J. P. "Van Fraassen on Propositional Attitudes." *Philosophical Studies* 42, no. 1 (1981): 101–110.
Reddam, J. P. Pragmatics and the Language of Belief. Ph.D. diss., University of Southern California, 1982.
Richard, M. "Temporalism and Eternalism." *Philosophical Studies* 39 (1981): 1–13.
Richard, M. "Tense, Propositions, and Meanings." *Philosophical Studies* 41 (1982): 337–351.
Richard, M. "Direct Reference and Ascriptions of Belief." *Journal of Philosophical Logic* 12 (1983): 425–452. Also in Salmon and Soames (forthcoming).
Russell, B. *Principles of Mathematics*. New York: Norton, 1903.
Russell, B. "On Denoting." *Mind* 14 (October 1905): 479–493. Also in Russell 1956 and Russell 1973.
Russell, B. "Knowledge by Acquaintance and Knowledge by Description." In Russell's *Mysticism and Logic and Other Essays* (London: Longmans, Green, 1911). Also in Salmon and Soames (forthcoming).
Russell, B. *The Problems of Philosophy*. New York: Oxford University Press, 1912.
Russell, B. "Descriptions." In Russell's *Introduction to Mathematical Philosophy* (London: Allen and Unwin, 1919).
Russell, B. *The Analysis of Mind*. London: Allen and Unwin, 1921.
Russell, B. "The Philosophy of Logical Atomism." In Russell 1956.
Russell, B. *Logic and Knowledge*, ed. R. C. Marsh. London: Allen and Unwin, 1956.
Russell, B. "Mr. Strawson on Referring." *Mind* 66 (July 1957): 385–389. Also in Russell 1973.
Russell, B. *Essays in Analysis*, ed. D. Lackey. London: Allen and Unwin, 1973.
Sainsbury, R. M. "On a Fregean Argument for the Distinctness of Sense and Reference." *Analysis* 43 (January 1983): 12–14.
Salmon, N. "Review of Linsky, 1977." *Journal of Philosophy* 76 (August 1979): 436–452.
Salmon, N. *Reference and Essence*. Princeton University Press, 1981.
Salmon, N. "Assertion and Incomplete Definite Descriptions." *Philosophical Studies* 42 (1982): 37–45.
Salmon, N. "Fregean Theory and the Four Worlds Paradox." *Philosophical Books* 25, no. 1 (1984): 7–11.
Salmon, N. "Tense and Information Content." In Salmon and Soames (forthcoming).
Salmon, N. "Reference and Information Content: Names and Descriptions." In *Handbook of Philosophical Logic IV*, ed. D. Gabbay and F. Guenthner (Dordrecht: Reidel, forthcoming).
Salmon, N., and S. Soames, eds. *Propositions and Attitudes*. Forthcoming.
Schiffer, S. "The Basis of Reference." *Erkenntnis* 13 (July 1978): 171–206.
Schiffer, S. "Indexicals and the Theory of Reference." *Synthèse* 49 (1981): 43–100.
Schwartz, S. *Naming, Necessity, and Natural Kinds*. Ithaca: Cornell University Press, 1977.
Searle, J. *Intentionality*. Cambridge University Press, 1983.
Sellars, W. "Putnam on Synonymity and Belief." *Analysis* 15 (April 1955): 117–120.
Sharvy, R. "Why a Class Can't Change Its Members." *Noûs* 2 (November 1968): 303–314.
Smullyan, A. F. "Modality and Descriptions." *Journal of Symbolic Logic* 13 (March 1948): 31–37. Also in Linsky 1971.

Soames, S. "Lost Innocence." *Linguistics and Philosophy* 8 (1985): 59–71.

Soames, S. "Direct Reference and Propositional Attitudes." In Salmon and Soames (forthcoming).

Sosa, E. "Propositional Attitudes *De Dicto* and *De Re*." *Journal of Philosophy* 71 (December 1975): 883–896.

Stalnaker, R. "Propositions." *Issues in the Philosophy of Language*, ed. A. MacKay and D. Merrill (New Haven: Yale University Press, 1976).

Stalnaker, R. "Assertion." In *Syntax and Semantics 9: Pragmatics*, ed. P. Cole (New York: Academic, 1978).

Stalnaker, R. "Indexical Belief." *Synthèse* 49, no. 1 (1981): 129–151.

Stalnaker, R. *Inquiry*. Cambridge, Mass.: MIT Press, 1984.

Strawson, P. F. "On Referring." *Mind* 59, no. 235 (1950): 320–344. Also in Strawson 1971.

Strawson, P. F. *Logico-Linguistic Papers*. London: Methuen, 1971.

Tienson, J. "Hesperus and Phosphorus." *Australasian Journal of Philosophy* 62, no. 1 (1984): 16–25.

Tye, M. "The Puzzle of Hesperus and Phosphorus." *Australasian Journal of Philosophy* 56, no. 3 (1978): 219–224.

Wettstein, H. "Indexical Reference and Propositional Content." *Philosophical Studies* 36 (1979): 91–100.

Wettstein, H. "Demonstrative Reference and Definite Descriptions." *Philosophical Studies* 40 (1981): 241–257.

Whitehead, A. N., and B. Russell. *Principia Mathematica*, second edition. Cambridge University Press, 1927.

Wiggins, D. "Identity-Statements." In *Analytical Philosophy*, second series, ed. R. J. Butler (Oxford: Blackwell, 1968).

Wiggins, D. "Frege's Problem of the Morning Star and the Evening Star." In *Studies on Frege II: Logic and the Philosophy of Language*, ed. M. Schirn (Stuttgart: Bad Canstatt, 1976).

Wittgenstein, L. *Tractatus Logico-Philosophicus*. London: Routledge and Kegan Paul, 1922.

Wittgenstein, L. *Philosophical Investigations*, tr. G. E. M. Anscombe. New York: Macmillan, 1953.

Wittgenstein, L. *Philosophical Grammar*. Berkeley: University of California Press, 1974.

Woodfield, A., ed. *Thought and Object*. New York: Oxford University Press, 1982.

Yagisawa, T. Belief and Synonymy. Ph.D. diss., Princeton University, 1981.

Yagisawa, T. "The Pseudo-Mates Argument." *Philosophical Review* 93, no. 3 (1984): 407–418.

Index of Theses

The Naive Theory of Information Value

Thesis I

(Declarative) sentences encode pieces of information, called *propositions*. The proposition encoded by a sentence, with respect to a given context, is its *information content* with respect to that context.

Thesis II

The information content, with respect to a given context, of a sentence is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the sentence, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices.

Thesis III

The information value (contribution to information content), with respect to a given context c, of any singular term is its referent with respect to c (and the time of c and the world of c).

Thesis IV

Any expression may be thought of as *referring*, with respect to a given context, time, and possible world, to its information value with respect to that context.

Thesis V

The information value, with respect to a given context, of an n-place first-order predicate is an n-place attribute (either a property or an n-ary relation)—ordinarily an attribute ascribed to the referents of the attached singular terms. Exceptions arise in connection with quotation marks and similar devices.

Thesis VI

The information value, with respect to a given context, of an n-adic sentential connective, is an attribute, ordinarily of the sorts of things that serve as referents for the operand sentences.

Thesis VII

The information value, with respect to a given context, of an n-adic quantifier or second-order predicate is an n-ary attribute, ordinarily of the sorts of things that serve as the referents for the operand first-order predicates.

Thesis VIII

The information value, with respect to a given context, of an operator other than a predicate, a connective, or a quantifier is an appropriate attribute (for sentence-forming operators) of, or operation (for other types of operators) on, the sorts of things that serve as referents for its appropriate operands.

Thesis IX

The information value, with respect to a given context, of a sentence is its information content, the encoded proposition.

The Singly Modified Naive Theory

Thesis I

(Declarative) sentences encode pieces of information, called *propositions*. The proposition encoded by a sentence, with respect to a given context, is its *information content* with respect to that context.

Thesis II

The information content, with respect to a given context, of a sentence is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the sentence, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices.

Thesis III′

The information value (contribution to information content), with respect to a given context *c*, of any simple singular term is its referent with respect to *c* (and the time of *c* and the world of *c*).

Thesis IV′

Any expression may be thought of as *referring*, with respect to a given context, time, and possible world, to its extension with respect to that context, time, and possible world.

Thesis V′

The information value, with respect to a given context, of a simple *n*-place first-order predicate is an *n*-place attribute (either a property or an *n*-ary relation)—ordinarily an attribute ascribed to the referents of the attached singular terms. Exceptions arise in connection with quotation marks and similar devices.

Thesis VI′

The information value, with respect to a given context, of a simple *n*-adic sentential connective is an attribute, ordinarily of the sorts of things that serve as referents for the operand sentences.

Thesis VII′

The information value, with respect to a given context, of a simple *n*-adic quantifier or second-order predicate is an *n*-ary attribute, ordinarily of the sorts of things that serve as referents for the operand first-order predicates.

Thesis VIII′

The information value, with respect to a given context, of a simple operator other than a predicate, a connective, or a quantifier is an appropriate attribute (for sentence-forming operators) or operation (for other types of operators), ordinarily an attribute of or an operation on the sorts of things that serve as referents for its appropriate operands.

Thesis IX′

The information value, with respect to a given context, of a typical compound expression, if any, is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the compound expression, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices, and may arise also in connection with compound predicates. The information value, with respect to a given context, of a sentence is its information content, the encoded proposition.

The (Doubly) Modified Naive Theory

Thesis I′

(Declarative) sentences encode pieces of information, called *propositions*. The proposition encoded by a sentence, with respect to a given context and time, is its *information content* with respect to that context and time.

Thesis II′

The information content, with respect to a given context and time, of a sentence is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the sentence, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices.

Thesis III″

The information value (contribution to information content), with respect to a given context *c* and time *t*, of any simple singular term is its referent with respect to *c* and *t* (and the world of *c*).

Thesis IV′

Any expression may be thought of as *referring*, with respect to a given context, time, and possible world, to its extension with respect to that context, time, and possible world.

Thesis V″

The information value, with respect to a given context *c* and time *t*, of a simple *n*-place first-order predicate is an *n*-place attribute, ordinarily an attribute temporally indexed to *t* (either a temporally indexed property or a temporally indexed *n*-ary relation)—ascribed to the referents of the attached singular terms. Exceptions arise in connection with quotation marks and similar devices.

Thesis VI″

The information value, with respect to a given context and time, of a simple *n*-adic sentential connective, is an attribute, ordinarily of the sorts of things that serve as referents for the operand sentences.

Thesis VII″

The information value, with respect to a given context *c* and time *t*, of a simple *n*-adic quantifier or second-order predicate is an *n*-ary attribute, ordinarily an attribute temporally indexed to *t* of the sorts of things that serve as referents for the operand first-order predicates.

Thesis VIII″

The information value, with respect to a given context and time, of a simple operator other than a predicate, a connective, or a quantifier is an appropriate attribute (for sentence-forming operators) or operation (for other types of operators), ordinarily an attribute of or an operation on the sorts of things that serve as referents for its appropriate operands.

Thesis IX″

The information value, with respect to a given context and time, of a typical compound expression, if any, is a complex, ordered entity (e.g. a sequence) whose constituents are semantically correlated systematically with expressions making up the compound expression, typically the simple (noncompound) component expressions. Exceptions arise in connection with quotation marks and similar devices, and may arise also in connection with compound predicates. The information value, with respect to a given context and time, of a sentence is its information content, the encoded proposition.

Index

Abstract entity, 25, 55
Acquaintance, 65, 67, 101, 109
  direct, 65
  with proposition, 107–108
Acquisition, of term, 70
Ambiguity, 98–99, 168
Analytic, 2, 133–135, 138, 166, 178
A posteriori, 2, 57, 78, 133, 137, 142, 176–178
  necessary, 141, 176
Appearance, 106–107, 110, 117
  visual, 73–74
Application, criteria of, 70, 82–83
Apprehension, mode of, 107
A priori, 2, 60, 69, 78, 133, 136, 141–142, 176–177
  with respect to context, 133
Argument from error, 69
Assent
  inward, 88–90, 110–111, 119, 130, 167
  verbal, 87, 90–91, 113, 167–171
Attitudes, de re, 20
Attribution
  de re, 8
  notional, 3
  propositional, 6, 26, 80, 153
  relational, 3, 174

Begriff, 160. *See also* Concept
Begriffsschrift (Frege), 11, 51, 79
BEL relation, 111–120, 126–127, 132, 171–173
Belief, 80–81, 88, 115, 171, 173–174
  de re, 19, 89, 102, 114, 169, 174
  withholding of, 101, 111–112, 120
Belief-predicate, 118
Belief-state, 173
Boër, S., 172
Bradley, R., 167
Burge, T., 171–172

Cappio, J., 158
Carnap, R., 138, 163
Castañeda, H., 8, 94, 129, 172
Causal chain, 70
Character, 14–15, 29, 31, 33, 70, 75, 165, 172
Chisholm, R., 172
Church, A., 22, 47, 138–140, 156–157, 159, 164
Cognition, 106, 167
Co-instantiation, 46
Common nouns, 161, 177, 179
Compositionality, 55–56, 61, 73–74, 77, 92, 157
Concept, 61, 160
  individual, 49, 125, 161
Conceptual association, 68
Conceptual content, 63, 65–69, 75, 124–125, 158, 166
  pure, 67, 75
Conceptual representation, 46, 63, 165
Content
  conceptual, 63, 65–69, 75, 124–125, 158, 166
  descriptive, 65
  semantic, 4, 19, 85, 117
Content base, 27, 29, 36, 39
Context, 14, 21, 28–38, 60, 64, 80, 154, 176
  doxastic, 81
  epistemic, 81
  opaque, 3
  temporal, 158
  transparent, 3
Contextual mechanism, 71

Contingent, 2, 178, 180
Contour, 15, 31
Contour operator, 39
Cordateness, 22

Demonstratives, 27, 46, 61-62, 70, 74, 125, 127, 171
Denotation, 64, 164. *See also* Reference
Denoting phrase, 45
Description
 definite, 12, 20–21, 28, 4[illegible], 52, 140, 161, 169–170
 noneternal, 42
Descriptional, 64
Designation, mode of, 51
Devitt, M., 165
Disquotation principle, 125, 132, 174
Donnellan, K., xi, 45, 78, 155, 180
Doxastic disposition, 110
Doxastic state, 112
Dummett, M., 80

Encoding relation, 68
Epistemological argument, 69
Erkenntniswerte, 1, 30, 47, 57, 91, 165
Essential property, 174
Eternalizing, 26
Eternal proposition, 27
Experts, of linguistic community, 70
Extension, 14, 22, 31, 33, 36, 53
Extensionality, principle of, 55
Extensional operator, 36, 38, 169

'Fido'-Fido theory, 16, 173
Fregean theory, 1, 6, 43, 120–121, 166
Frege, G., 1, 11–12, 19–20, 25, 29–30, 35, 46–54, 63–64, 77, 107, 160, 162, 164
Frege's Law, 57, 60, 73–74, 77, 80, 162
Frege's Puzzle, 6–7, 11, 19, 55, 57, 62, 77, 79, 91–92, 127
 new, 87–92, 101
Frege's Strategy, 73
 generalized, 73–76, 166

Gedanke, 1, 29, 47, 59
Gettier, E., 82
Gödel, K., 22, 157
Guise, 104–105, 109, 113, 115, 117, 127, 164

Hacking, I., 153

Higher-order property, 156
Hobbes, T., 51

Identity, 12, 51
Identity sentence, 49, 51
Identity statement, 51–52, 58, 74
Indexical, 14, 25, 28, 35, 42, 63–64, 66, 70, 74, 125, 165, 173, 176
Indexicality, 34–37
Information, 12, 55, 59, 120, 154
 eternalness of, 24–27
 pragmatically imparted, 58, 60, 69, 78–79, 84–85, 92, 100, 114, 127, 137, 141, 167
 semantically encoded, 59–60, 72, 78–79, 84–85, 91–92, 100, 114, 127, 137, 141, 163, 167, 173
Informational equivalence, 90
Information content, 1, 7, 13, 17, 23–24, 29, 34, 36, 41, 48, 57–58, 60, 91, 107, 116, 120, 122, 130, 153, 157, 163, 169, 174
Information-content base, 27
Information value, 14, 16–18, 21, 23–25, 27–36, 41–42, 55, 61, 63–72, 74, 124, 140, 157, 161, 166
 conceptual theory of, 63–70
 contextual theory of, 70–71
 verbal theory of, 71–72
Information-value base, 27–29, 32–33, 40
Informativeness, 47, 57, 60–61, 73–78, 91, 163, 167
Instantiation, 46
Intension, 13–14, 31, 33, 38, 153
Intensional entity, 49, 63, 170
Intensional operator, 3, 38
Intensional semantics, 165

Kamp, H., 158
Kaplan, D., xi, 1, 7, 14, 16, 31, 45, 65, 153, 156, 159–163, 170, 173, 180
Kneale, M., 26, 157
Kneale, W., 26, 157
Knowledge, 80–82
 analysis of, 83
 extension of, 11, 57, 60
Kripke, S., 2, 7, 45, 65, 72, 77, 94, 131, 137, 139, 161, 164–165, 170, 174
Kripke's Puzzle, 129–132

Language, ordinary, 82–83. *See also* Usage

Leibniz's Law, 57, 77, 80
Lewis, D., 168, 172, 175
Linguistic activity, 53–54, 162
Linguistic chain, 70–71
Linguistic community, 70
Linguistic convention, 52, 84, 162
Linguistic entity, 61
Linsky, L., 48, 161
Locational operator, 35
Logical attributes, 131, 174
Logical equivalence, 22
Logical truth, 60, 135, 174–175
Logical validity, 133, 176
Lycan, W., 172

McKay, T., 167
Marcus, R. B., 45, 174
Material equivalence, 116
Mates, B., 7, 138, 179
Meaning, 13, 31–32, 60, 135, 138, 160, 165, 173, 176
Meaningfulness, 47
Measurement term, 139
Metaphysical necessity, 138. *See also* Truth, necessary
Mill, J. S., 51, 77, 161–162
Modal argument, 69, 140
Modal equivalence, 116
Modality, de re, 19–20
Modal operator, 38, 153
Montague, R., 159
Moore, G. E., 153

Naive theory, 11, 16–17, 19, 23, 25, 32, 50, 162
 modified, 19, 24, 31, 41, 43, 53–54, 61, 63, 70, 73, 75, 79, 84, 92–93, 98, 111, 114, 119–120, 126, 128, 135, 167, 171, 176
 of perception, 122–123
Names
 nonreferring, 7, 127, 168–169
 proper, 21, 43, 54, 61–62, 70–73, 86, 126 (*see also* Single-word term)
Naming and Necessity (Kripke), xii, 66, 135
Natural-kind term, 72
Negative existential, 127, 168–169

Orthodox theory, 64–66, 119–122, 125–126

Perception
 analysis of, 83
 de re, 122–123
 naive and sophisticated theories of, 122–123
Perry, J., 7, 45, 172–173
Possible world, 13, 33–34, 63–64, 154–155, 168
Possible-world semantics, 2, 153
Predicative component, 57
Presentation, mode of, 47, 49, 52, 64, 120, 160
Price, H. H., 167
Prior, A. N., 158
Program, 28, 31–35, 134
Proposition, 1, 16–18, 22–27, 41–42, 78, 81, 113, 130–131, 155, 172
 diagonal, 168
 general, 1, 6, 161
 open, 155
 singular, 6, 8, 20, 30, 46, 49, 61, 89, 91, 93, 99, 102, 108–109, 112–116, 124, 130, 153, 161, 168–169, 171, 174
Propositional attitude, 6, 26, 80, 153
Propositional function, 28, 45–46, 156, 160
Propositional term, 4
Proposition matrix, 24, 27–28, 39, 155
Psychology, 120, 124, 173–174
Purely conceptual content, 67, 75
Purely qualitative concepts, 66
Putnam, H., 65–66, 72, 78, 138, 163, 165

Quine, W. V. O., 94, 129, 134, 156, 170, 174
Quotation-disquotation principle, 132

Recognition failure, 103–109, 115, 118, 126–127
Reddam, P., 167
Reference, 2, 7, 21, 51, 64, 161, 180
 anaphoric, 6
 direct, 50, 65, 69, 78, 81, 161, 165, 172
Referent, 11, 18, 47, 70–75, 181
Referring, 18, 23, 41
Reinateness, 22
Relationally descriptive term, 64
Relativity theory, 50
Representation, 47
Representation content, 52, 166
Richard, M., 26, 158

Russell, B., 7–8, 19, 45–46, 63–64, 77, 107, 158–161, 164
Ryle, G., 16

Sainsbury, R. M., 167
Schedule, 28, 31–33, 39
Schiffer, S., 173
Semantic attribute, 68
Semantic content, 4, 19, 85, 117 (*see also* Information value)
Semantic mechanism, 63
Semantics
 formal, 4–5
 intensional, 165
 possible-world, 2, 153
 temporal, 36
Semantic significance, 58
Semantic stipulation, 53, 162
Semantic value, 4, 15, 26, 28, 31, 33–34, 63, 153
Sense, fixing of, 53
Sense (Sinn), 4, 11, 30, 47, 51–52, 65, 120–121, 125–126
Sentential connective, 17–18, 23, 37, 41–42
Set-abstraction term, 168–169
Sharvy, R., 162
Significance, cognitive, 79, 165
Single-word term, 72–73, 127. *See also* Name, proper
Soames, S., 158
Speech act, 20
Stalnaker, R., 168
Stipulation, 53–54, 138, 140–141, 162
 semantic, 53, 162
Substitutivity, 80, 86–87
Superintension, 28, 32–33, 39
Superintensional operator, 39
Swartz, N., 167
Synonymy, 163, 176–177
Synthetic, 2, 135

Tarski, A., 4
Temporal operator, 35–39, 43
Temporally indexed attribute, 40
Thoroughly descriptive term, 64
Thought, 1–2, 25, 29, 47, 49, 60, 107, 120, 174
 de re, 3–4 (*see also* Attribution, relational)
 language of, 120
Tractatus Logico-Philosophicus (Wittgenstein), 11
Translation, 138–139, 170
Truth
 logical, 60, 135, 174–175
 necessary, 69, 176 (*see also* Metaphysical necessity)
Truth-condition, 19, 37, 65–66
Truth-functional, 16, 36
Truth-value, 13–14, 25–26, 37, 63, 99, 117
Twin-earth argument, 66, 124, 172–175
Tye, M., 167

Über Sinn und Bedeutung, 52
Understanding, 87, 175–176
Usage, 54, 81–84, 162
Utterance, 59, 154

Variable, 4–5, 7, 64

Wiggins, D., 153, 161
Wittgenstein, L., 11–12, 153

Yagisawa, T., 167