

From Brain to Cosmos

Preliminary Revised Edition

Mark F. Sharlow

From Brain to Cosmos: Preliminary Revised Edition

Copyright © 2013 Mark F. Sharlow
All rights reserved

The first edition of *From Brain to Cosmos* was published in 2001 by Universal Publishers/uPUBLISH.com and is copyright © 2001 by Mark F. Sharlow.

The items “Notes on From Brain to Cosmos: Questions and Answers about Subjective Fact” and “How Subjective Fact Ties Language to Reality” were available on the author's website before 2013. The second of these items bears its own copyright notice. Parts of the first edition of *From Brain to Cosmos* also were available on the author's website before 2013.

Introduction to the New Edition

It has been a while since I wrote the first edition of *From Brain to Cosmos*. The book you are now reading is a revised and expanded edition designed for release as an e-book. I have cut and pasted material from the first edition together with some new material to get a book that I believe will be clearer and more informative than the original book.

In this introduction I'll tell you what the book is about. Then I'll give you a brief guide to what's in the book, followed by a longer chapter-by-chapter guide. Along the way, I'll clarify some points that might be confusing for a new reader of the book.

My main purpose in writing *From Brain to Cosmos* was not radical or even all that original. I was only trying to explore a question that has been of interest to many philosophers, including Descartes, Russell, Carnap, and Husserl. Here is the question:

How much can we learn about the real world by tracing the logical consequences of facts about how things appear?

From Brain to Cosmos

I stated this question right up front, on page 1 of chapter 1. I stated it in slightly different words there, but it's the same question.

In the book I did not try to find a complete answer to this question. That would have been much too ambitious a goal. Instead, I proposed a few partial answers to special cases of the question. Some of these partial answers turned out to be interesting enough to justify the effort spent in finding them.

Starting from these partial answers, I was able to suggest a possible answer to another old question: what is the relationship between mind and physical reality? I did not try to solve the mind-brain problem, but I proposed a new view of the relationship between physical reality and the conscious subject who observes it. This view does not fit comfortably into the usual categories of materialism, dualism or idealism. It could be best described as a modest form of idealism fully compatible with the scientific view of mind. According to this view, the facts of the physical world have a kind of logical dependence (not causal dependence) on facts about experience – but this dependence is fully compatible with a materialistic explanation of mind. I realize that any suggestion that smacks of idealism is likely to raise eyebrows nowadays, but the “idealism” proposed here is fully compatible with the scientific outlook.

Aside from this blending of materialism and idealism, I also was able to suggest some new ideas about the flow of time and our experience of that flow, and about the temporal structure of conscious subjects. Among other points, I argued that there can be an objectively real flow of time even if there are no objective tenses, and that

From Brain to Cosmos

conscious observers are (in one sense) temporally extended entities.

These are a few of the main ideas that grew out of the studies in this book. The details of these ideas, and the arguments for them, are in the chapters of the book.

A Quick Guide to the Book

Here is a brief guide to the book. After this brief guide, I will lay out a longer, chapter-by-chapter guide to what's in the book.

In Chapters 1 through 3, I develop the concepts and logical tools that I will use throughout the book. These include the ideas of *subjective fact* and of *instance of seeming*, and a few other ideas defined in terms of them.

In the rest of the book, I use these concepts and tools to analyze a variety of philosophical issues and problems. (I do not try to solve most of these problems. Usually I just try to clarify them and to propose some solutions in special cases.)

In Chapters 4 through 6, I use these concepts and tools to analyze a few philosophical issues about mind and knowledge.

In Chapters 7 through 10, I use the same set of concepts and tools to analyze some philosophical questions about time and our awareness of time.

From Brain to Cosmos

In Chapters 11 through 14, I use the same concepts and tools again to study some further problems about mind and persons.

There is no Chapter 15 in this book. I deleted it because it was out of date, but I left the original chapter numbering the same as in the first edition.

In Chapter 16 I make a few closing remarks.

After Chapter 16 come three appendices, two of which were not in the original book. These might be of interest if you have read most of the chapters.

The next several pages contain the detailed chapter-by-chapter guide. This guide contains advice about what parts to read first and what to leave for later.

Contents

This table of contents is the same as the one in the first edition, except for one cross-out. I have not rebuilt it for this edition. The front matter has changed since then, and Chapter 15 was deleted.

Chapter 1

In this chapter I state the question that is the book's theme. Then I discuss the question in detail, both to clarify its meaning and to show why it is of interest. Also,

From Brain to Cosmos

I state exactly what I plan to do in the rest of the book.

This chapter (after the first question) contains a lot of details. I included these because I wanted to prevent confusion about what I am doing in the book. Put simply, I am *not* trying to find secure foundations for all human knowledge, and I am *not* trying to rebut absolute skepticism as Descartes tried to do. (I do not think such a rebuttal is necessary.) I also mention other overly ambitious things that I am not trying to do. I am only trying to explore the logical consequences of facts about appearance, and to find out what we can learn about objective reality in this way.

If you find Chapter 1 too long-winded, you can skim through the later parts quickly on a first reading – but try not to pass judgment on what I am doing until you have read all of this chapter and found out about all the silly things I am *not* trying to do.

Chapter 2

In this chapter I analyze the notion of *seeming* – an idea that plays the starring role in the book's main question. I explore some logical features of facts about how things seem and of facts about what seems to be the case. I introduce the notion of an “instance of seeming,” which plays a crucial role later in the book. Also, I introduce the kindred notions of “subjective fact” and “subjective being,” which also figure heavily in the rest of the book. (Naturalists, don't panic; subjective being is not at all what it sounds like.) I clear up some logical niceties about these three notions – such as the semantics of quantification over instances of seeming. A large part of

From Brain to Cosmos

this chapter is devoted to examples of seeming or appearance. These examples are meant to flesh out the definitions of the notions introduced in the chapter.

Late in the chapter, I discuss the relationship of all these ideas to some issues in the philosophy of mind. I point out that I am not proposing a theory of consciousness here. Also, I discuss the relationship of the idea of subjective fact to relativism and absolutism with regard to truth.

Chapter 3

In this chapter I develop the idea of an “instance of seeming” further than I did in Chapter 2. I show that these instances seem like events in some respects, and I adopt the more picturesque term “consciousness event” as a replacement for “instance of seeming.” (This choice of words may have been a bad move on my part, since a “consciousness event” is not necessarily an event at all. Just keep in mind that a consciousness event is actually just an instance of seeming.)

In the rest of the chapter I study the notion of consciousness event in gory detail. You can skip a lot of the details and examples on a first reading, but be prepared to come back to them if you are puzzled by anything I say in later chapters.

I finish with some discussion of the logical properties of consciousness events, from the standpoint of modal logic and possible-worlds semantics. This part too can be skipped or skimmed on a first reading.

From Brain to Cosmos

The Next Part of the Book

In Chapters 4 through 15, I take up some classic philosophical problems and try to analyze them using the tools I developed in Chapters 1-3.

Warning: I want to emphasize that my aim in Chapters 4-15 is not just to study some traditional philosophical problems, but to study those problems *using the tools developed earlier in the book*. For this reason, my studies of these problems sometimes look very different from what one usually reads in journal papers on these problems. I am well aware of this difference. I am not ignoring the way analytic philosophers usually do things. I am just trying to approach these problems in a specific, novel way. Be patient with me, and you will find in the end that I have not strayed that far from the kind of philosophy you are familiar with.

Chapter 4

In this chapter I begin the book's project in earnest. I try to show, using the concepts from Chapters 2 and 3, that reliable knowledge of one's own *immediately past experience* is possible. This seemingly trivial result is the first substantive consequence that I try to deduce from facts about how things seem.

Chapter 5

In this chapter I attack a weightier philosophical problem: the well-known problem of personal identity

From Brain to Cosmos

through time. I don't try to solve this problem as a whole. Instead, I develop a sketch for a theory of streams of consciousness, using the concepts from Chapters 2 through 4. (A note to philosophers of mind: Don't worry, these "streams" don't have to be truly continuous.) Also, I develop an account of subjective temporal succession. These partial results shed light on the problem of personal identity, even though I don't propose a theory of personal identity as such.

Chapter 6

In this chapter I use the ideas from earlier chapters to approach the problem of knowledge of other minds. I model this knowledge using the logic of consciousness events (recall Chapter 3). I argue that acquaintance with other subjects' mental states is possible even if some aspects of consciousness turn out to be truly first-person and private.

Chapters 7, 8 and 9 deal with philosophical issues about time. You can probably skip these chapters on a first reading if you aren't interested in the philosophy of time. (You can go back to them later if anything I say about time seems puzzling.)

Chapter 7

In this chapter I take a look at an old topic in the philosophy of science: the flow of time. First I examine the views of one philosopher who took the flow of time to be ultimately real. I contrast this viewpoint with other

From Brain to Cosmos

ideas that portray time's flow as merely psychological. I propose my own definition of temporal flux, which does not fully capture the intuitive notion of time's flow but does capture some important features of that notion. Then I examine the notion of happening (as it applies to events) and the relationship between tense and temporal flux. I conclude that there can be real temporal flux, becoming, and happening even in a world in which there are *not* objectively real tenses. Also, I conclude that any stream of consciousness (as defined in Chapter 5) must involve temporal flux.

Chapters 8 and 9

In these chapters I continue the discussion of time that I started in Chapter 7. In Chapter 8 I argue that perceived events can have a kind of tense even if there is no objectively real tense in the world. In Chapter 9 I show how the existence of real happening might be reconciled with the tenseless existence of events in a spacetime.

Chapter 10

In this chapter I examine another problem related to time: how objects persist through time. I argue for a view of physical objects and conscious subjects as temporally extended items. My view, unlike some similar views, allows for the full reality of temporal flux and happening. Finally, I point out a way to defeat Humean doubts about the existence of a persisting self. (This final section is of interest if one takes such doubts seriously.)

From Brain to Cosmos

Chapter 11

In this chapter I use some of the ideas in Chapters 1-5 to analyze a few concepts from psychology and the philosophy of mind. I take up the topics of subconscious mental life and the disunity of the conscious self. I suggest that so-called “unconscious” mental contents might actually be *conscious* mental contents of a sort, and that these contents might seem unconscious only because they are difficult for the conscious subject to know about. Also, I argue that a conscious subject can be a genuine individual even if there is a lot of disunity within the subject.

Chapter 12

In this chapter I examine the problem of personal identity through time. This is a continuation of what I did in Chapter 5. Instead of taking on the problem of personal identity directly, I study the related concept of the identity of *conscious subjects* through time. (This kind of identity might not always coincide with personal identity, but it is easier to analyze within the framework of Chapters 1-5.) I focus on the often-discussed problems of dividing conscious subjects, including standard puzzles about split-brain operations.

Chapter 13

In this chapter I attack a small piece of the very big problem of the relationship between mind and matter. I do *not* propose a solution to the mind-brain problem here. I

From Brain to Cosmos

only try to explore the relationship between physical existence and experience, in much the same way that Mill, Russell, and others have done. I start with a brief discussion of the notion of truth. In most of the rest of the chapter I develop a new view of the relationship between experience and the existence of the physical world. At the end of the chapter I say a few words about the currently fashionable forms of relativism.

Chapter 13 is the most venturesome chapter in the book. I want to emphasize that the ideas I propose here are offered more as possibilities than as firm conclusions.

Chapter 14

In this chapter I use the apparatus developed in Chapters 1-3 to approach the question in the chapter's title: "Which beings are conscious?" The question of which living organisms are conscious is an old question in philosophy. The question of whether computers can be conscious is a modern variant of that question. I develop a new conceptual framework for studying these questions, though in the end I do not propose definitive answers to either question.

Chapter 15

The first edition of the book had a Chapter 15, in which I commented on some issues in the philosophy of religion. I have done more work in this area since then, and this chapter became obsolete. Hence I left Chapter 15 out of this edition. Readers interested in my work in this area should track down my other books and papers that

From Brain to Cosmos

deal with the subject.

Chapter 16

This chapter is a postscript to the book. It doesn't tell you much unless you have read the book.

Notes and Works Cited

The reference notes here are from the first edition and pertain only to the material from that edition. Since I did not update that material (except to move or delete parts), I have not tried to update the references. The Works Cited section also is from the first edition.

After the End of the Book

After the end of the book, I added a few appendices that expand on the material in the book.

Appendix A: Notes on *From Brain to Cosmos*. This addresses a few questions and objections that might occur to the reader.

Appendix B: Preface to the First Edition. This was the preface to the original published version of the book. I now think it belongs at the end of the book. It addresses some possible doubts that might arise during a reading of the later chapters of the book.

Appendix C: How Subjective Fact Ties Language to Reality. This paper might be of interest to those worried

From Brain to Cosmos

about postmodern relativisms.

Closing Remarks

This concludes my introduction to the revised edition of *From Brain to Cosmos*. If you have any questions, feel free to contact me. As of the time of this writing, my e-mail is:

msharlow@usermail.com

If my e-mail address ever changes, you might be able to find my new contact information on my website, or in my profile at some of the philosophy preprint archives where my papers are archived. As of this writing, my website is at:

<http://www.eskimo.com/~msharlow>

I'd like to thank the many people who contributed to this book in various ways, including the early readers of the note sets that led to this book, the publisher of the first edition, and readers of that edition. I'd also like to thank my professors and fellow students from my days as a philosophy student. They started me thinking about some of the issues I discuss in the book.

- Mark Sharlow
2013

From Brain to Cosmos

Cautionary Note on Page Numbering and Contents

As I said earlier, I constructed this e-book by cutting and pasting on the first edition. I have not yet tried to impose a new page numbering on the file. As a result, the page numbering jumps around a bit. Also, the table of contents given here is from the first edition (except for the cross-outs). Since then, items have been deleted or added, and the page numbering has changed. This book is more like a preprint than a finished product – but I'm not going to promise that I will or won't neaten it up someday.

Contents

Preface	xv
Acknowledgments	xxi
Chapter 1. On How Things Seem to You	1
The Roots of Knowledge	
A Method with a Past	
The Agenda of This Book	
What This Book Is Not About	
Appendix: What Makes This Project New	

The Many Meanings of "Seem"

Instances of Seeming

A Technical Note: Quantifying Over
Instances

Subjective Fact

Facts and Abstract Objects:
A Cautionary Note

Subjective Fact Revisited

Subjective Being

More Examples of Subjective Fact

Consciousness and How Things Seem

An Aside on Theories of Consciousness

An Aside for Philosophers of Mind

Subjective Fact and Intentionality

Chapter 2 (continued)

The Subjective and the Objective

Relativism and Subjectivity

Subjective Fact as a Fundamental
Notion

Chapter 3. **The Happenings Within** 53

Consciousness Events

Consciousness Events: Some Examples

Consciousness Events and Experience

The Timing of Consciousness Events

Unusual Consciousness Events: Some
Possibilities

The Logic of Consciousness Events:
Three Appendices

*Appendix A: Subjective Fact and
Modality*

Chapter 3 (continued)

*Appendix B: Consciousness Events
as Incomplete Worlds*

*Appendix C: Are Consciousness
Events Logically Consistent?*

Chapter 4. **On Knowing What Just Happened** 69

Certain Knowledge About Experience

Seeming to Seem

The Redundancy of Seeming

The PSR and Knowledge

Chapter 5. **Conscious Beings and Their Histories** 89

Personal Identity: An Introduction

An Agenda

The Vagueness of Personal Identity

Continuance and Subjective Duration

The Stream of Consciousness

Chapter 5 (continued)

An Empty Objection Defeated

Subject Identity During Periods of
Unconsciousness

Three Technical Notes

*Note 1: How To Formalize Subjective
Time*

*Note 2: The Ontology of Subject
Histories*

*Note 3: The Topology of Subjective
Time*

Chapter 6. **Knowledge of Other Minds** 123

The Privacy of the Psychological

Perception of Other Subjects'
Subjective Content

The Logic of Noticing Other Subjects'
Mental States

Chapter 6 (continued)

The Fallibility of Perception of Other
Minds

Summing Up

Chapter 7. **The Flow of Time** 139

Temporal Flux and Bergsonian
Duration

Temporal Flux in Sense Experience

How Events Happen

Some Technical Notes on Instantaneous
Events

More Technical Notes: Duration and
Happening

Tense, Temporal Motion, and Time's
Flow

Re-Charting the River of Time

The Objective Reality of Time's Flow

Chapter 7 (continued)

From Consciousness to Objective
Time

Chapter 8. **The Experience of Time** 177

Subjective Tenses

Subjective Tense and Subjective
Ontology

Temporal Modalities: A Technical
Note

Chapter 9. **Spacetime and Happening** 191

The Timeless Happening of Events

Global Happening and Super-time

Subjective Flux Is an Aspect of
Tenseless Happening

Minkowski versus Bergson: A Peace
Proposal

Chapter 10. **Conscious Beings and Physical Things** 203

What Are Diachronic Objects?

What Is a Persisting Object?

Why Things Appear to Be Temporally
Unextended

Closing Remarks on the Self

Chapter 11. **The Structure of the Self** 223

The Conscious Subject: A Review

A Subject Is Not an Ego

Unknown Experiences

The Nature of the Unconscious:
An Hypothesis

The Absolute Unity of the Subject

The Logic of Confusion: Some
Technical Points

Chapter 11 (continued)

The Self: Real Unity in Apparent
Multiplicity

Chapter 12. **Personal Identity: Some Problems** 245

Splitting Subjects

Could I Be Two People?

Splitting Subjects with Psychological
Compartmentation

Splitting Subjects Again, in Less
Detail

Commissurotomy and Subject
Splitting

Merging Subjects

In Conclusion

Chapter 13. **Mind and Matter** 273

The Notion of Objective Truth

Chapter 13 (continued)

From Appearance to Knowledge

The Experiential Principle: What
It Is Not

Reality and Subjective Facts

Idealism

Idealism and Physicalism

Recursive Idealism

Appendix: An Epistle from the
Far Side of Relativism

Chapter 14. **Which Beings Are Conscious?** 315

Immediate Knowledge About
Consciousness

Inferential Knowledge About
Consciousness

Consciousness Is a Sharp Quality

The Continuous Alteration Principle

Chapter 14 (continued)

An Objection Rebutted

The Shrinking Mind Revisited

Consciousness in the Biological World

So Which Beings Are Conscious?

Chapter 15. **God, Time and Eternity** 339

Science versus God: A Nonproblem

God the Supreme Subject

Greenwich and God's Wristwatch

Is Theism a Thesis About Value?

Time and Eternal Life

What Makes Experience Real?

What Makes Experience Real? (The Brief
Version)

Immortality and the Brain: A Speculation

Chapter 16. Postscript	377
Notes	383
Works Cited	419

Chapter 1

On How Things Seem to You

This book is a study of some long-standing philosophical puzzles about the nature of the universe in which we live. These puzzles have to do with two of the most intriguing features of that universe: the ones we call consciousness and time. In this book, I will approach these philosophical riddles in an unusual way: by attempting to answer one far-reaching question which will shed some light upon all of them. While trying to answer this question, I will arrive at conclusions about consciousness and time which will hint at a new view of the nature of reality itself.

Because I do not want to keep the reader waiting, I will state this question now. The precise meaning of this question — and its vital importance — will become clearer over the course of the next few pages.

What can we learn about the nature of reality — about what really exists — by deducing the consequences of facts about how things seem to conscious beings?

It may not yet be clear how this question relates to the topics of time and of consciousness. To illustrate how the question links up with these topics, I will begin by offering some observations about human knowledge. These observations suggest that we can learn much about both time and consciousness by considering the consequences of facts about how the world seems to us.

The Roots of Knowledge

All that we know, we know by means of mental or psychological activities. These activities, which include such thinking, perception, and insight among others, always involve conscious mental states in some manner. To see, feel, or otherwise sense an object, you must be conscious. To think about a fact or to gain insight into a problem, you must be conscious. Thus, consciousness seems to play a central role in the process of knowing.

Any alleged knowledge which is acquired and used without the involvement of any conscious processes is not knowledge in the customary sense of that word. A completely unconscious machine does not truly know anything, even if it processes items of information (like the facts of arithmetic) which a conscious human would regard as "knowledge." If someone attributes knowledge to a completely unconscious system, I would argue that that person understands the meaning of the word "knowledge" differently from the rest of us. Knowledge, as ordinarily

understood, cannot exist in the complete absence of consciousness.

In view of the centrality of consciousness to knowledge, one expects the ideas of consciousness and of experience to play important roles in any attempt to investigate the nature of knowledge.

The central feature of consciousness — the feature that makes conscious beings truly conscious — is the fact that for a conscious being, there is a *way things seem*. There is a way that things appear to that being — a way that its world seems to it to be. In contrast, a being for which things do not seem to be any way at all is not conscious, at least not in any standard sense of that word.¹ The possession of a way things seem is one of the key features, and perhaps the defining feature, of consciousness.

The idea that consciousness is the possession of a way things seem is equivalent, at least in its essentials, to the ways that a number of philosophers have characterized consciousness. For our purposes, the most relevant of these characterizations is due to Thomas Nagel, who pointed out that an organism's being conscious involves the organism's having "something it is like to *be* that organism."² Later I will discuss some other connections between philosophers' views of consciousness and the notion of having a way things seem.

When you want to learn about your own consciousness, the easiest way to do so is to pay attention to how things seem to you. To realize this, consider the following question: How do you know you are conscious right now?

From Brain to Cosmos

Answering this question is easy: you just look around you, or listen to the background noise, or think a thought — and you can tell at once that you are conscious. Things *seem* a certain way to you now; your external environment looks a certain way or sounds a certain way, and the internal act of thinking feels a certain way. You know that you are conscious because there is a way things and events *seem* to you. If you became unconscious right now, the way things seem to you would vanish. In brief: the fact that you are conscious follows from the fact that things seem a certain way to you. It follows from a fact about how things seem.

This example shows that the idea of learning about consciousness by studying how things seem is not a new or radical idea. In hindsight, the possibility of such learning is almost too obvious to be mentioned. Aside from learning about consciousness, we may ask whether it is possible to learn about the *external* world by studying the way things seem. The possibility of this kind of knowledge also is obvious: indeed, all the knowledge we have was attained with the help of information about how things seem. All of us rely constantly upon our experiences, and hence upon information about how things seem, to help us navigate through the day. Scientific knowledge also requires facts about how things seem; scientists base their conclusions largely upon observations, and hence indirectly upon facts about how things appear, or seem, to conscious observers.

A Method with a Past

The method of learning about reality by studying the way things seem has turned out to be quite familiar. It is little wonder that philosophers have tried repeatedly to extend this method of learning so as to create new strategies for the acquisition of knowledge. Some noted philosophers have tried to uncover new knowledge about external reality by constructing *deductive arguments* which begin from premises about how things seem. These thinkers have tried to obtain some knowledge by examining the *necessary consequences* of facts about how things seem, instead of merely drawing conclusions from their experiences in the customary ways, either scientific or commonsensical.

The first philosopher to conspicuously attempt a project of this sort was René Descartes. In the 1600's, Descartes proposed a method which he thought could reduce the amount of error in human knowledge. Today this method is known as *systematic doubt*. Here I will merely summarize the highlights of this method, as it is set forth in Descartes' main work, *The Meditations Concerning First Philosophy*.³

The method of systematic doubt requires its user to perform a breathtaking exercise of the imagination. To use this method, you pretend that the world around you might be an illusion, and that everything you think you know might be false. Then you search for reasons which show that your knowledge is correct — reasons which do not depend on

assuming that *any* of the things around you might be real. If any part of your knowledge can be justified through reasons this persuasive, then that part should be free of error. (It is extremely important to note that using Descartes' method is not the same as *actually doubting* the existence of the external world. Descartes was not rejecting reality; he was only playing a game of sorts to test the soundness of his knowledge.)

Through this method, Descartes arrived at what he believed to be an absolutely certain truth: the fact of his own existence. In effect, Descartes pointed out that even if the external world were an illusion, there would have to be *someone to have the illusion*. Descartes argued that experience, and particularly thinking, cannot occur unless there is a self or mind to undergo the experience. He tried to show, in effect, that experience cannot occur unless a mind or self also exists. This amounts to a claim that a fact about reality (the existence of a self) can be inferred from facts about experience — which are almost the same as facts about how things seem.

In his quest to pass from experience to reality, Descartes did not stop at the proof of his own existence (which most of us feel was too obvious to require proof). Descartes went on to construct weightier arguments — arguments in support of major philosophical conclusions about mind, matter, and the cosmos. Thus, Descartes tried to deduce some truly philosophical knowledge — knowledge about the nature of reality — from facts about experience.

Later philosophers have cast a great deal of doubt upon

Descartes' program. The most interesting criticism is one which casts doubt upon Descartes' belief that conscious experience requires the existence of a conscious self. In the eighteenth century, David Hume raised the very strange possibility that conscious experience might exist *without* a self. Hume argued that the existence of a persisting self is an illusion produced by "successive perceptions."⁴ In the twentieth century, Bertrand Russell argued that the existence of a self does not really follow from the existence of experience. In effect, Russell pointed out that it is safer to use the fact that *experience is occurring* as a starting point, instead of the fact of one's own existence.⁵

Am "I" an entity (physical or spiritual) that continues from moment to moment, or is there only a series of experiences which includes the feeling that "I" persist through time? Hume and Russell realized that this question cannot be answered with a glib "Of *course* I persist!" Such a common-sense "answer" does nothing to answer the question, since this answer could be blurted out and sincerely felt to be true by a conscious brain, even if that brain contained no persisting self beyond the stream of experience! The answer to Hume's question is even less obvious now, in this age of artificial intelligence and neuroscience, than it was in Hume's time. But the question of whether there really is a persisting self is not the point here; I will take up that question much later in the book. The point is that the human mind can at least conceive of the possibility of conscious experience without a persisting self. And if we can at least imagine this possibility, then we

cannot accept Descartes' argument for the existence of the self without thinking twice. We cannot pass directly from the fact that thinking or experience is occurring to the conclusion that a thinking or experiencing self, and not just the current scrap of experience or thought, really exists. As Russell pointed out, it is safer to start from the bare fact that *experience is occurring*. And since the most easily knowable part of experience is the way things seem, this amounts to a challenge to derive human knowledge from facts about how things seem.

In the twentieth century, the quest for a deductive path from experience to world remained alive and well. The analytic philosophers Bertrand Russell (whom I just mentioned) and Rudolf Carnap both tried to derive large parts of human knowledge by using the technique of *logical construction*.⁶ This technique, which I will not discuss in detail here, requires one to begin with a class of items which one takes as real and fundamental. Both Russell and Carnap placed experiences among their fundamental entities. Carnap used "elementary experiences", or experiences as such.⁷ Russell used "particulars," among which he included "sense-data,"⁸ which may be regarded, more or less, as fragments of experiences.⁹ Thus, two of the most influential analytic philosophers of the twentieth century tried, in effect, to deduce much of human knowledge from facts about experience.

Another twentieth-century movement whose adherents tried to derive knowledge rather directly from experience

was the phenomenological movement. This movement, which is quite large and active today, grew out of the phenomenological techniques which Edmund Husserl introduced for the study of experience. Phenomenology is too large a subject to be summarized in a few sentences. Here I will confine myself to some very general (and somewhat oversimplified) remarks, and will concentrate on phenomenology as conceived by Husserl, which is not the only brand of phenomenology on the market today.¹⁰

Phenomenology, as Husserl conceived it, begins with an attempt to investigate one's experiences in a special manner. Such investigation begins with a mental maneuver often called "epoché," which begins with an attempt to observe one's surroundings without viewing them as having objective reality beyond one's experience. In the frame of mind which this procedure creates, one can carry out investigations of experience itself — of the world taken merely as experience, instead of as a real, external world. Phenomenology does not claim that the world is unreal; it merely makes use, in its investigations, of an attitude in which the world is at first taken purely as an object of experience.

Phenomenology comes close to being an attempt to derive facts about reality from facts about how things seem. However, it does not quite amount to such an attempt, since its primary concern is not with logical deductions from facts about how things seem, but rather with the exploration of experience itself. Also, phenomenologists tend to experiment with, or manipulate, their experiences in ways

which go far beyond the simple noting of how things seem to oneself. The procedure of describing how things seem and making deductions from the resulting descriptions can be applied to any kind of experience, phenomenological or not.

Any philosophical project which relies heavily on the description of how things seem will have some points of contact with phenomenological thought. I will point out one of these in Chapter 6 when I discuss the problem of other minds. But overall, this book does not take a phenomenological approach to its subject matter.

The Agenda of This Book

The preceding discussion shows that the project I am undertaking in this book is neither new nor strange. At heart, it is simply a new variation on a centuries-old theme which has been taken seriously by such philosophers as Descartes, Husserl, Russell and Carnap. However, my approach to this theme will be different from the approaches taken by these philosophers. I will now point out some of the details of my approach which make it novel. The difference between my project and those of the philosophers just mentioned will be discussed in more detail in the Appendix to this chapter.

In this book, I will attempt to find out how much knowledge about the real world can be deduced from facts about how things seem. I will ask what, if anything, one can

learn about what really exists — about what philosophers call the *ontology* of the world — simply by making deductions from mere facts about *what seems to be the case*. In other words, I will attempt to find out how much we can learn about the world simply by tracing the logical consequences of facts about what conscious observers experience.

This task, however, is but the first part of the project. While carrying out this task, I also will formulate certain philosophical theses which are suggested or motivated by the conclusions drawn from facts about how things seem. These theses will not follow rigorously from facts about how things seem. Rather, the conclusions drawn from such facts will render the philosophical theses plausible, or at least make them more intuitively appealing.

Thus, my overall line of argument will consist of two parts. First, there will be a core of deductive arguments which begin with facts about how things seem. Second, there will be a set of less rigorous arguments for various philosophical views; these views are motivated by the conclusions of the first, more rigorous, set of arguments. Some of the conclusions of the second set of arguments will be but slight extensions of familiar philosophical ideas. Others will be more venturesome.

The amount of philosophy which one can do in this way will turn out to be surprisingly large. During our deductive voyage from experience to world, I will present arguments about several traditional philosophical problems, including those of personal identity, temporal flux, and the realism-

idealism debate. In some cases I will be able to devise new solutions to these problems, or to defend solutions which have been proposed by others. All of these proposed solutions will, of course, be partial and tentative; they are not meant to be conclusive. Nevertheless, they have far-reaching implications when taken together. This will become evident toward the end of the book, where I will argue that these partial and tentative solutions push us toward a new overall view of reality.

The view of reality which I finally will propose is somewhat different from any of the standard pictures of the nature of reality. This view does not fall without resistance into the customary categories of realism, idealism, dualism, and the like. Indeed, this conception of reality has more in common with some current ideas in theoretical physics than with any of the traditional ontological standpoints. This new view affirms, with modern physics, that the world around us is not the kind of thing that we normally feel it to be. It also suggests that the relationships between mind and matter and between the subjective and objective standpoints are relationships of a surprising and intricate sort. But I am getting ahead of myself. I will describe this new view of reality in detail in later chapters of the book.

The above statement of purpose is intended to let the reader know roughly what the book is about. The details and the conclusions will be laid out in the chapters to come.

What This Book Is Not About

Now that I have said what I am going to do, I wish to point out some things which I am *not* going to do.

First, I am not attempting to rebut skepticism, as Descartes tried to do. I am not going to ask whether absolute skepticism is a real threat to knowledge, nor will I raise the question of the usefulness of Descartes' method of systematic doubt. Of course, if one feels that a rebuttal to absolute skepticism is necessary, then one should be intensely interested in the question which I am trying to answer here. One may even find one's skepticism rebutted before one gets to the notes. However, one can think that the problem of skepticism is a nonproblem and still find the book as interesting as a confirmed doubter would find it.

Aside from the traditional problem of skepticism and knowledge, there is another issue about skepticism which this book does address. This is the issue of relativism. In recent years there has been a flood of attempts to show that all our knowledge is relative in some sense. Such relativism can be understood and evaluated within the framework of the present project. I will have more to say about this possibility later.

Second, I am *not* trying to provide secure foundations for human knowledge, as Descartes, Carnap and others wished to do. My attempt to deduce some knowledge from facts about how things seem might lead someone to try to label me as a foundationalist. This label would be inaccurate for

two reasons: I am not claiming that *all* knowledge can be deduced from facts about how things seem, and I hold that such facts are grounded in still other facts. This book is neutral with regard to foundationalism.

Third, I will not address the general problem of the reliability of knowledge obtained through experience. My arguments will not tell us under what conditions such knowledge is reliable, except in some very limited cases. At most, those arguments will agree with the commonsense conviction that the objects around us are real — while also showing that none of those objects is quite what it seems to be.

Appendix: What Makes This Project New

The project of this book is indebted to the work of the philosophers mentioned in the preceding sections. However, my undertaking differs from these earlier projects in a number of ways. The first difference, which I already have mentioned, is in my starting point. Descartes began his project with assumptions about minds and mental contents. Hume, Russell and Carnap founded their theories on the assumed existence of mental objects of various sorts — "impressions" and "ideas" for Hume, "sense-data" for Russell, and "elementary experiences" for Carnap.¹¹ All of these mental objects were supposed to be real items. I will begin further back, without assuming the existence either of minds or of mental objects, but instead assuming that it

sometimes seems as if this or that is the case. My goal is not to derive conclusions from an assumption that mental items exist — even if that assumption appears to someone to be self-evident. Rather, my aim is to derive some knowledge from facts about how things seem. That is a different goal.

This difference in starting point makes the present project immune to certain criticisms which bedevil theories like those of Russell, Carnap and Hume. A philosophical strategy which begins with the assumption that intramental objects are real is vulnerable to criticism on the grounds that such objects do not exist. Recent philosophers of mind even have attempted to deny that feelings and thoughts really exist at all.¹² It does not matter whether one thinks that this denial is true or is simply silly; we still must not beg the question of its correctness at this stage of our project. By using facts about how things seem as a starting point for our project, we can dodge most criticisms from this direction. No philosopher of mind ever can successfully maintain that things do not sometimes *seem* one way or another. Philosophical or scientific theories about the mind may provide accounts of what it is for things to seem a certain way, but they cannot do away with the fact that things do seem a certain way to you now. A recent example of a theory which denies many of the putative subjective features of consciousness is the materialistic theory proposed by Daniel C. Dennett.¹³ Yet Dennett has virtually admitted, in his book on that theory, that his preferred method for the study of consciousness (what he calls "heterophenomenology") does not attempt to defeat claims

about how things seem.¹⁴

My project also differs in many other particulars from the earlier projects mentioned above. Unlike Descartes, I will not adopt total skepticism as part of my method; indeed, my overall line of argument will be quite different from that of Descartes. Unlike Russell and Carnap, I will not try to find logical constructions which can substitute for concrete objects. Unlike phenomenologists, I will not make use of any special mental acts directed at experience. Instead, I simply will analyze facts about how things seem.

Chapter 2

Into the Subjective World

The project of this book, as described in the last chapter, requires us to focus upon one particular feature of consciousness. This feature is the one we call the way things seem. In this chapter, I will introduce a set of concepts for the description of the ways things seem to conscious observers. These concepts will let us begin the first part of this book's project — the deductive transition from experience to world.

The Many Meanings of "Seem"

Before beginning my remarks on the way things seem, I should clarify what I mean by the phrase "the way things seem," and in particular by the word "seem." At this stage of my project, I have no need for rigorous definitions of these terms, nor am I able to supply such definitions in any noncircular manner. Instead, I will assume that the reader is familiar with the ordinary usage of the verb "to seem."

From Brain to Cosmos

The meaning of "to seem" requires some clarification, for in ordinary usage this verb has several distinct senses. I will use "to seem" in one specific sense: the sense which embodies what one ordinarily calls awareness or experience. Here are some examples of uses of "seem" which reflect this usage of the word:

If a square green object drifts into your field of vision, then it *seems* to you that there is a square green object.

If you just heard a brief loud noise, then it *seems* to you that a loud noise occurred.

If I feel cold, then it *seems* to me that it is cold.

If you are thinking about philosophy, and suddenly realize that you are thinking, then it *seems* to you that you are thinking.

If I have just finished turning around rapidly for a few minutes, then it *seems* to me that my surroundings are turning — even though I know that they really are not.

If I look out the window at a gray, darkening sky, and I get a vaguely ominous feeling from this sight, then the gray sky seems somehow ominous to me — even if I know that it really is not a threat to me at all.

If I have a hallucination of a green dragon, then it *seems*

to me that there is a green dragon — even though there really is no green dragon.

In all of these examples, the word "seem" is used to indicate what someone's surroundings are experienced to be like, or (what is the same) to the appearance or feel of things. We might call this sense of "to seem" the *consciousness sense* — the meaning of "to seem" which you use when you talk about how your world appears or feels to you.¹ The consciousness sense of "to seem" is not the only meaning which this verb can have. "To seem" often is used in a way which expresses *belief* rather than consciousness or experience. For example, if I say "It seems to me that George will win the election," I probably mean much the same thing as when I say "I believe that George will win the election." (Perhaps I also am trying to indicate that this belief is rather tentative.) In this case, "It seems to me that George will win the election" does not mean that I really am *experiencing* George's future electoral victory. I do not have *any* sort of experience of this victory, for the victory still is in the future (and may never even happen). George's victory, or the fact that George wins at some future time, simply is not part of what the world appears like to me now.

One can think of other uses of "to seem" which point to meanings other than that of awareness or experience as such. I will get back to some of these uses later. For our present purposes, the only important sense of "to seem" is the consciousness sense.

By marking off the consciousness sense of "to seem"

from other senses, I am not suggesting that the consciousness sense of "to seem" is free of vagueness or ambiguity. Later in this book, I will bring up an argument from the philosophy of mind which suggests that the consciousness sense of "to seem" contains some irreducible vagueness. But this argument, even if right, is no threat to the project of this book. My characterization of this consciousness sense of "seem" is precise enough for my present purposes.

Throughout the book, when I use the phrase "facts about how things seem" and its variants, I will mean facts about how things seem, with "seem" given its consciousness sense.

Note that my use of the word "things" in "how things seem" is not meant to restrict the subject matter to the appearances of *things*, like tables, chairs or stones. My intended meaning is more general: "how things seem to you" means how your surroundings in general seem to you, or (to borrow a phrase of Nagel's) "what it is like *for*" you.² Perhaps the more colloquial expression "how it seems" would be less misleading than "how things seem," though even the former locution is somewhat misleading (there is no "it," or entity, which seems that way). But I will stick with "how things seem" for want of a better phrase.

Instances of Seeming

For any conscious being, there are facts about how things seem to that being. Things may seem different ways to a conscious being at different times. The existence of a way things seem is not the only interesting or important feature of consciousness, but it is the one which will most concern us in this book. Our project requires us to think about how things seem to various conscious observers, and to use this information to try to find out something about the nature of reality.

The project which we are undertaking requires us to construct arguments whose premises include facts about how things seem. To do this, we must be able to express, in some suitable language, various facts about how things seem. At first sight this appears easy to do: to state a fact about how things seem, one simply takes a statement and prefixes the phrase "It seems that" (or some equivalent thereof) to the statement. For example: "It seems that there is a dog coming through the door." "It seems that the sky is blue."

Unfortunately, statements like these leave out some essential information about the way things seem. This omission occurs because of a peculiar property of facts about how things seem — a property which I will now explore.³

Suppose that I am trying to describe my experience, and I utter the sentence "It seems like there is some red there."

Later on, I am trying to describe a new experience, and I utter the sentence "It doesn't seem like there is any red there." Of course, these two sentences do not contradict one another, as they might appear to do if taken out of context. Why not? The answer, of course, is that the two sentences were uttered at two different times. Or, one might say instead, because the two sentences pertained to two different *instances*, or examples, of seeming: the appearance of my surroundings to me in one instance, and the appearance of my surroundings to me later on.

Now suppose that I again utter, in all sincerity, "It seems like there is some red there." Simultaneously, you utter "It doesn't seem like there is any red there." Do these two statements describe a contradictory situation? Obviously not. The two sentences were uttered by different conscious subjects, and hence pertain to two different points of view. Or, one might say instead, the two sentences pertain to different *instances*, or examples, of seeming: the appearance of things to me now, and the appearance of things to you now.

These two excessively simple examples illustrate a fundamental logical property of seeming which is almost too obvious to notice: namely, that *seeming has instances*. Things do not merely seem such-and-such a way; rather, they seem such-and-such a way in this or that *instance*, or *occurrence*, of seeming. My perceiving red, my failing to perceive red, and your failing to perceive red occurred in three distinct instances of seeming. In one of these instances, it was the case that there seems to be some red. In

the other two instances, this was not the case. In general, there are many different occasions on which things seem to be certain ways. The way things seem on one of these occasions does not have to cohere with the way things seem on another such occasion.

To construct adequate descriptions of how things seem, we need to be able to talk about how things seem in different instances. It is not enough merely to say "It seems that P," where P is some statement. If we only say this much, then we have not said all we can say about the instance of seeming in question. In either of the "red" examples above, if we tried to describe how things seem without worrying about which instance was involved, we would get something like this: "It seems that there is some red there. It does not seem that there is some red there." Since the situations described were not contradictory, it is clear that something is missing from the description.

It is possible, of course, to fill out descriptions of how things seem by indexing them with an observer and a time: "It seems to me now that there is some red there." "It seems to Henry at 12 noon that there is a dog coming in the door." This familiar device allows us to specify in what instance things seem a certain way. However, these sentences have a serious flaw: they are not simply sentences about how things seem. Instead, they also convey information about how things really are, apart from how things seem. When I say "It seems to Henry at noon that there is a dog coming in the door," I am saying that a conscious subject (Henry) has a certain kind of subjective experience at a certain time

(noon). Such a sentence cannot be true unless the subject exists at the stated time: if there is no Henry, then it cannot seem to Henry that there is a dog coming in the door, and if there is no time called "noon," then nothing can seem like anything to anyone at noon. Thus, the sentence "It seems to Henry at noon that there is a dog coming in the door" cannot be true unless there is a conscious subject and a time — provided that we take that sentence literally. We may know for certain that Henry exists and that there is such a time as noon, but we do not want to assert these facts when we are trying to assert a sentence that expresses *only* a fact about how things seem.

A better way to bring instances of seeming into our language is simply to talk about the instances of seeming themselves. Suppose that I want to say that things seem a certain way to me now. The way things seem to me now is the way things seem in a certain instance. I can refer to this instance of seeming — or, at least, I can perform a linguistic act which appears, at first glance, to be one of reference, specifically of reference to an instance of seeming. I can assign the instance of seeming a symbol, say "x," which I use to refer to that instance — or at least to perform the linguistic act of apparent reference which I just mentioned. Hence instead of saying something like "It seems to me now that P," I could say this: "In instance x, it seems that P." This sentence is intelligible regardless of whether or not "x" actually refers.

At first glance, this trick might appear to involve reference to real items — entities called "instances of

seeming" — and hence to go beyond the description of how things seem. A little reflection shows that this is not a problem, for two reasons. First, it is possible to read the sentence "In instance x, it seems that P" in such a way that "x" does not refer. I will have more to say about this possibility later. Second, even if one thinks that "x" does refer, such reference does not commit one to anything beyond the facts about how things seem. An instance of seeming is nothing more than a particular occasion or example of how things seem. Hence to say that there is an instance of seeming is to say nothing more than that things seem some way. When we assert that there are instances of seeming, we really are asserting nothing more than we do when we assert that things seem some way or other. We are not asserting the existence of some extra objects called "instances of seeming" in addition to the facts about how things seem.⁴ An instance of seeming is not a separate entity, above and beyond all facts about how things seem. Rather, it is a feature of those facts. We can talk about instances of seeming, give them "names," and even reason about them without assuming anything worse than the view that things can seem different ways on different occasions. (Whether we legitimately can quantify over instances of seeming is a question I will take up later.)

Aside from the above remarks and some notes later on, I will not discuss the semantics of the symbols used to refer to instances of seeming. The semantics of such symbols must be very similar to the semantics of natural-language expressions like "the instance of seeming in which it seemed

that the clock struck 12" — expressions which, on their face, appear to refer to single instances of seeming. Since these expressions make sense, it follows that artificial expressions which perform the same function (such as the "x" introduced two paragraphs ago) make sense also. A philosopher of language might try to pin down the semantics further, perhaps asking first whether the relation between an expression of this sort and an instance of seeming is really an example of reference. Since such questions are not vital to my project, I will avoid them here. To try to preserve neutrality on such questions, I will call the expressions in question *tags* instead of names.

It is characteristic of facts about how things seem that any such fact belongs to an instance of seeming. At least this is true of all such facts of which I know. For our present purposes, we are safe in assuming that it is true for all facts about how things seem, period. The following argument shows why. Suppose — just for the sake of argument — that there were some facts about how things seem which did not belong to any instance of seeming. Then we could invent one fictitious instance of seeming for each such fact, and say that each such fact belongs to its own instance of seeming. For fact P, say, we introduce a term " x_P ," and stipulate that "P seems to be the case in instance x_P " is true. All we are doing is adding several new symbols to our language, and deciding to use them in a certain way. Having done this, we can ask whether there really are instance of seeming of which the new symbols, " x_P " and the like, are

tags. We quickly find that we can make a case for the existence of such instances, given that our other tags (the ones we had before we invented the new symbols) really are tags for instances of seeming. The sentence "there is an instance of seeming called ' x_P '" tells us nothing over and above the claim that certain facts about how things seem are the case — specifically, that it is the case that it seems that P. Thus, the claim that " x_P " is a tag of an instance of seeming is a claim of the same sort as a claim that any other putative tag of an instance of seeming really is such a tag. We cannot be wrong in making such a claim if the appropriate facts about how things seem (in this case, just the fact P) really are the case, and if also we are using the putative tag in a certain way.

It is important to note that more than one fact may seem to be the case in the same instance of seeming. In a particular instance, it may seem that there is something pink and it also may seem that there is an elephant.

A Technical Note: Quantifying Over Instances

Earlier I mentioned the issue of quantification over instances of seeming. In this section I will show that one can quantify over instances of seeming without making any existential commitments other than those involved in asserting facts about how things seem. Some of my earlier statements may look more plausible in view of the conclusions of this section. Readers not deeply interested in

issues of existential commitment can skip this section without much loss of continuity.

Suppose that you want to assert that it seems that P. You know that it does not simply and unqualifiedly seem that P; instead, it seems that P in *this* instance. To state this last fact, you can invent a new phrase, "It x-seems that." Take "It x-seems that P" to mean that it seems that P in *this* particular instance — that is, in the instance in which you found it seemed that P. Now change "It x-seems that P" stylistically, to read "In instance x, it seems that P."

To quantify over instances of seeming, one can quantify over the x in sentences like this, giving the quantifiers their substitutional readings. The use of the substitutional reading here does not have to be defended on the grounds that the objectual reading would involve us in unwanted existential commitments.⁵ Rather, the substitutional reading of the quantifiers simply captures the intuitive idea of the existence of an instance of seeming better than does the objectual reading. An instance of seeming, I have said, is not an entity existing over and above the facts about how things seem. The sole condition for the existence of an instance of seeming is that things seem a certain way. If we give the quantifier the substitutional reading in "There exists an instance of seeming x such that...", then claims that instances of seeming exist will boil down to claims that certain facts about seeming are the case. This is precisely the outcome we want, since to say that there exists an instance of seeming is to say that things seem some way or other and nothing more.

Subjective Fact

Now that the concept of instance of seeming is in place, I can introduce another general concept which will play an important role throughout this book. This is the concept of *subjective fact*.

We have seen that any fact about how things seem is a fact about how things seem in a particular instance. Suppose that P is a statement describing how things seem, without reference to the instance of seeming involved. (To continue our earlier example, if Henry sees a dog come in the door, then P might be "A dog comes in the door.") Let x be a tag for an instance of seeming. Then we will say that the sentence

In instance x, it seems that P

is a *subjective-fact sentence*. We will call the fact expressed by this sentence — namely, the fact that in instance x, it seems that P — a *subjective fact*.

Facts and Abstract Objects: A Cautionary Note

By speaking here of facts and of subjective facts, I do not mean to commit myself to the existence of facts as abstract objects. Some philosophers⁶ have held that facts (or states of affairs) are among the real constituents of the world.

When I speak of facts, I am using that word in a prephilosophical sense: the same sense that a journalist or a physicist uses when speaking of "the facts." I do not believe that any of my talk about facts commits me to the existence of facts as abstract objects, for reasons detailed in the next two paragraphs.

It is possible to talk of "facts" without taking sides on the question of the reality of abstract objects. As a simple example, one can truthfully say "It is a fact that Fido is a dog" without believing in any abstract objects at all. In this example, the reference to facts clearly is redundant; the sentence is simply a paraphrase of "Fido is a dog." But in a less trivial example, one can say "The fact that Fido is a dog is a fact about Fido" without presupposing the existence of any abstract objects. One can, if one wishes, paraphrase away the reference to facts in this sentence in favor of a more nominalistic reference to sentences. But whether or not one uses such a paraphrase, one can simply read the sentence itself in a way which avoids commitments to facts. Dennett has pointed out⁷ that, although it is possible to talk sensibly about people's voices, the meaningfulness of such talk does not entail the existence of a special entity called a voice, above and beyond items like a person's vocal organs, sounds, and so forth. Read in this way, talk about voices becomes somewhat figurative. One can understand talk about *facts* in much the same way. It is possible for a journalist to have the facts about a traffic accident or an election, even if there are no abstract entities called facts.

When I speak of facts in this book — for example, when I

assert that there are facts about how things seem — my talk about facts may be understood in this noncommittal way. I am not claiming that there are facts *qua* abstract objects, or that there are not; this question is beyond the scope of this book. If the reader prefers, he or she can construe all my fact-talk realistically, and interpret it as discourse about entities called facts. But I am going to speak freely about facts without insisting on this interpretation, and will leave the question of the reality of facts wide open.

My decision to ignore the question of the ontological status of facts does *not* mean that I think this question is unimportant. I have made this decision because an answer to this question would have little or no bearing on the things I wish to do in this book.

The preceding remarks about fact-talk also are applicable, with appropriate changes, to my use of words which seem to denote abstract objects, such as "set," "class," and "property." I will use these words, but the reader is free to interpret them without the help of abstract objects if he or she so wishes. In this book, I do not wish to make any claims about the reality of abstract objects; hence any statement along the lines of "there is a property such that..." may be given either a nominalistic or a realistic reading at the reader's discretion. Hopefully all of my abstract-object statements can be handled this way. I think the problem of the ontological status of abstract objects is important, but I am not going to take it up in this book.

Subjective Fact Revisited

At times, I will use sentences of forms other than "In instance x, it seems that P" to express subjective facts. Often I will say instead that *P is the case for x*. This form is only a stylistic variant of "In instance x, it seems that P," yet it has a certain psychological advantage: it highlights an important characteristic of subjective facts. This characteristic is a parallel between the idea of *seeming* and that of *truth*. If something seems to an observer to be the case, then that something plays the role of a fact or truth in the observer's subjective "world." Such a fact embodies (in part) what the world is like for that subject — what is true *for that subject's awareness*, regardless of what (if anything) is true in reality.

Sometimes I will stretch this terminology even further, by saying that P is the case for S, where S is the conscious subject to whom the instance of seeming X belongs. This terminology will become more useful when I discuss conscious subjects in detail. So far, I have not explored the relationship between instances of seeming and conscious subjects. However, since it is intuitively clear that conscious subjects "have" instances of seeming, we can make intuitive sense of the locution "P is the case for subject S."

In a similar vein, when I speak of conscious subjects, I may sometimes speak of them as "having" instances of seeming, or I may speak of instances of seeming as "being

in" subjects. These ways of speaking have fairly clear intuitive meanings, even though I have not yet analyzed the notion of a conscious subject. To a conscious subject, things seem certain ways; hence that subject is associated in a certain manner with instances of seeming. Later on, I will develop an account of the relationship between a conscious subject and its instances of seeming.

All of the new locutions which I have introduced in this section are, at bottom, simply new ways of describing how things seem. The concept of subjective fact provides us with an idiom for the description of how things seem. All of these other locutions, such as "is the case for" and the like, are simply variations on this idiom, hopefully more intuitively appealing than the unadorned language of subjective fact.

Subjective Being

Another terminology which I will use on occasion is that of *subjective being*. Subjective being is a special case of subjective fact; it is defined in terms of subjective fact in the following manner.

Consider once again the locution "In instance x, it seems that P." In some cases, P will be a positive existential statement — a statement saying that there are objects of some sort. In this case, we will say that objects of that sort *exist for* x. (The notion of something being or existing "for" something is not new; neither are expressions like "exist for"

and "being for." These words and notions can be found in the writings of phenomenologists and Hegelians, but I am not committing myself to these lines of thought by using such words.⁸) Here is an example of the use of "exists for": If P is "There is an orange square," then we may say "An orange square exists for x." This latter sentence is just an equivalent of "In instance x, it seems that there is an orange square." As with subjective fact, I sometimes will stretch this terminology to make it applicable to subjects as well as to instances of seeming. Specifically, I may say that something exists for a *subject*, instead of for an instance of seeming which that subject has.

Like the idiom of subjective fact, the idiom of subjective being highlights the parallel between seeming and truth. If it seems to an observer that a thing of some kind exists, then a thing of that kind plays the role of an object or entity in the observer's subjective "world," even if there really are no things of that kind. Occasionally, when I need to distinguish real being from subjective being, I will refer to real being as objective being.

It is obvious, but still important to note, that "objects" which exist for an instance of seeming need not really exist at all. The fact that pink elephants exist for an instance of seeming does not imply that there are any pink elephants. Some philosophers have suggested that there are nonexistent objects as well as existent objects.⁹ The claim that pink elephants exist for an instance of seeming emphatically does not mean that there are pink elephants which are objects of any sort — existent or nonexistent. It only means that there

seem, in that instance, to be pink elephants. Whether there are nonexistent pink elephants is a question which I will leave open.

In many cases, when a fact is the case for an instance of seeming, this implies that there is something which exists for that instance. For example, if for instance *x* there is an orange square, then an orange square exists for *x*. However, it is possible for something to be the case for an instance of seeming, without anything existing *for* that instance. For an example of this, consider how it feels to be in a moderately hot place. When one is in such a place, there is a definite way that one's surroundings seem — yet this does not consist in an experience of any particular object. Moods provide still other examples of subjective fact without subjective being. When one is in a mood, the world just seems a certain way; there is no apparent object in whose perception the mood consists.

More Examples of Subjective Fact

Earlier in this chapter, I gave several examples to flesh out the consciousness sense of the verb "to seem." These were, in effect, examples of subjective fact. In this section I will present more examples of subjective fact. These examples will point up a number of interesting properties (some trivial, some not) of subjective fact.

In most of these examples, a situation is described in the language of seeming, then redescribed in the idiom of

subjective fact. The new description is parenthesized. In each sentence, the symbol "x" is a tag for the instance of seeming to which the subjective fact belongs.

Two ordinary examples of subjective fact:

(1) It seems to me now that there is a rectangular thing in front of me. (In the instance x, it seems that there is a rectangular thing in front of me.)

(2) It seems to me that there is a letter *e* in front of me. (In the instance x, it seems that there is a letter *e* in front of me.)

Example (1) illustrates an obvious principle: the fact that such-and-such is the case for someone does *not* imply that such-and-such *actually* is the case. (The screen in front of me actually is not rectangular, it only looks rectangular.)

Example (2) illustrates a different principle: that subjective fact can involve complex perceptions as well as supposedly "simple" sensations. The facts which can be the case for you are not restricted to very simple sensory facts. It can seem to you now that there is a letter *e*, although perceiving a particular pattern of black and white patches as an *e* is not a simple process from the standpoint of neurophysiology. If you merely saw a black-and-white pattern which you did not recognize as a letter, and then reasoned to the conclusion that that pattern was an *e*, then you would know there is an *e*, but it would not *be the case*

for you that there is an *e*.

The prime examples of things that are the case for someone without being the case are facts that seem to be true during hallucinations.

(3) It seems to me now that there is a huge green dragon. (In the instance *x*, it seems that there is a huge green dragon.)

Perceptual illusions also yield many examples of this sort.

(4) It seems to me that there is a square thing. (In the instance *x*, it seems that there is a square thing. Or: In the instance *x*, it seems that something is square.)

(In actual fact, the thing is not quite square and is tilted. But if I don't look carefully, I miss this fact, and the thing seems square.)

Facts about past events can seem to be the case.

(5) It seems to me that I went to work this morning. (In the instance *x*, it seems that I went to work this morning.)

The "seeming" described in example (5) occurs when I try to recall this morning's happenings. To arrive at the conclusion that I went to work, I do not have to infer consciously that I went to work; it simply seems to me that I

did so. Thus, something can seem to be the case when something is *remembered*. Facts about the past can seem to be the case in this way. When I remember today's commute, it seems to me that I went to work. (Of course, it does not seem to me that I *am going* to work.)

Occasionally, near future facts can seem to be the case, though most cases of anticipation of the future do not involve future facts seeming to be the case.

(6) It seems to me that I am about to catch a basketball. (In the instance x, it seems that I am about to catch a basketball.)

(This sentence can be true a split second before one actually grasps the ball — when the ball is on one's fingers and feels "caught." Of course, in reality, the ball could go either way; the ball's being caught is not assured. There is no precognition at work here.)

Just as subjective fact must never be confused with truth, so also it must never be confused with belief. In the following example, assume that I am nowhere near my desk, and that I believe that there is a pile of papers on my desk.

(7) It does not seem to me that there is a pile of papers on my desk. However, it also does not seem to me that there is *not* a pile of papers on my desk. (It is not the case that in the instance x, it seems that there is a pile of papers on my desk. Nor is it the case that in the

instance x, it seems that there is not a pile of papers on my desk.)

In the preceding example, my desk and its contents, as they are *now*, do not "seem" to me at all.

Example (7) illustrates my earlier cautionary remarks about the use of "seem" to express belief. English speakers often use "It seems to me that..." as a synonym of "I think that...", "I believe that...", "I opine that...", or "I suspect that..." Examples of such usages: "It seems to me that George will win the election." "It seems to me that this equation is right." "It seems to me that our party is right about this." "It seems to me that the suspect is guilty." These sentences express senses of "to seem" different from the consciousness sense. A similar cautionary remark should be made for alternative uses of "to appear." (I may exploit these alternative uses myself in this book; hopefully the distinction will be clear from the context.)

When a fact is too distant from me to be part of my "inner world" at all, then it neither is nor is not the case for me. Example (7) illustrates this. So do the next two examples.

(8) It does not seem to me that Pluto is directly overhead. (It is not the case that in the instance x, it seems that Pluto is directly overhead.)

(9) However, it does not seem to me that Pluto is *not* directly overhead. (It is not the case that in the

instance x, it seems that Pluto is not directly overhead.)

In these two examples, there simply is nothing in my experience right now which gives me any indication of the position of Pluto.

Examples (7), (8) and (9) should make it clear that "In the instance x, it seems that P" and "In the instance x, it seems that not-P" may both be false at once. This point will become important later.

Facts about events that are partly present and partly past to you can seem to be the case for you.

(10) It seems to me that this is the scariest part of the movie so far. (In the instance x, it seems that this is the scariest part of the movie so far.)

These ten examples should help to make more concrete the sense of "to seem" which the idiom of subjective fact is supposed to capture. To be the case for you is to seem to you to be the case — in precisely this sense of "seem."

In passing, I should mention that one can use sentences which express subjective facts to describe how things seem to a subject, even if that subject cannot describe his experiences in words. If Henry is aphasic but it seems to Henry now that there is a flash of green light, then it seems that way to Henry even if Henry cannot find the words to describe what he saw. Even if Henry is not aphasic, but has an experience so overwhelming or unique that he cannot put it into words (and you, if you had a similar experience, could

not put it into words either), you still can describe this experience within the language of subjective fact. Simply adopt a symbol, say "Q," and define it to mean "things are the way they seemed to Henry during his ineffable experience." Then it will be the case, for a certain instance of seeming x, that "In instance x, it seems that Q."

Consciousness and How Things Seem

The consciousness sense of "to seem" is closely related to several ideas about consciousness put forward by philosophers of mind. I will remark very briefly on some of these ideas.

The most central feature of consciousness — the feature that makes conscious beings genuinely conscious — is the fact that to a conscious being, things seem to be some way or other. For a nonconscious entity, things cannot seem to be any way at all. As I pointed out in Chapter 1, a being for which things do not seem to be any way at all is not genuinely conscious. The kind of seeming which makes conscious beings truly conscious is embodied in the consciousness sense of the verb "to seem."

Various philosophers' characterizations of consciousness appear to agree in their essential features with the one in the preceding paragraph. The most helpful of these characterizations, for our purposes, is the one due to Thomas Nagel. Nagel has suggested that the distinguishing characteristic of a conscious organism is the existence of

"something that it is like to *be* that organism — something it is like *for* the organism."¹⁰ Nagel also has suggested that consciousness involves the possession of a "point of view."¹¹ I would add this: we cannot say that a being has a point of view in any relevant sense unless there is a way that things seem, in the consciousness sense, to that being.

John R. Searle has emphasized the importance of the "first-person" aspects of consciousness.¹² But we can safely say that to have first-person character, consciousness must involve the having of a special perspective on one's surroundings, or on the world. That is, there must be a way that things seem.

A. J. Ayer once contrasted the experiencing of actions " 'from the inside'" and " 'from the outside'";¹³ the former perspective is that of the subject who is performing the actions. This distinction also implies that a conscious being has a special point of view on its own actions — that the way it relates to those actions is somehow fundamentally different from the way in which other beings relate to them.

One need not accept all of the philosophical views of Nagel, Searle, Ayer, or any other particular philosopher to find these philosophers' descriptions of consciousness useful and illuminating. Even Dennett, whose view of consciousness¹⁴ is very reductionistic, declines to allow his proposed method for the study of the mind to defeat conscious beings' claims about the way things seem to them.¹⁵

The above mentioned ideas about insides or viewpoints

all appear to capture, more or less accurately, the same essential intuition about what consciousness is. Consciousness is having a way that things seem for oneself. To be conscious is to have a subjective realm of apparent facts and things which seem to be true or real. This picture of consciousness is what emerges when we search for the lowest common denominator of the ideas just discussed: those of having an "inside," having a "point of view," having Nagel's "something it is like *for*" oneself (discussed above), and so on.

There may be alternative readings of the word "conscious" on which things need not seem one way or another for a conscious being.¹⁶ But these readings, if they exist at all, are not the ones we usually have in mind when we say things like "I am conscious." Ordinarily, we would not consider a zombie — that is, a hypothetical being with humanlike behaviors but without a subjective life — as conscious.¹⁷ If such beings are possible, and you were to become such a being, you would go unconscious. In this book I will not worry about any other meanings of the word "conscious." For my purposes, a conscious being is a being for which things seem to be some way or other.

An Aside on Theories of Consciousness

Before leaving the topic of philosophical theories of mind, I want to avert a possible misunderstanding about the contents of this chapter. I wish to emphasize that the claims

I am making in this chapter do not comprise a "theory of consciousness" in any sense of that term.

Philosophical and scientific theories of the nature of consciousness are intended to inform us about what sort of a phenomenon consciousness really is, or to make some lesser positive assertions about the nature of consciousness. Such theories lead to different views of the nature of claims about how things seem or appear. Dennett's theory, for instance, portrays claims about how things seem as mere "narratives" produced by zombie-like machines.¹⁸ Descartes' dualistic theory implies that such claims are accurate descriptions of states of a nonphysical soul. Most theories of consciousness lie somewhere between these two extremes.

The conclusions to which I have come in this chapter do not commit us to any particular theory of consciousness. They are neutral among the existing theories of consciousness. One can use the concepts introduced here — those of instance of seeming, of subjective fact, and of subjective being — to help one state facts about how things seem, while systematically ignoring the question of what ultimately makes those facts true.

Throughout most of this book I will try to answer certain questions about consciousness while avoiding the question of what consciousness really is. Much later — and as a result of having done this systematic ignoring — I will be in a position to venture some conclusions about the nature of consciousness. But even then, I will not attempt to offer a complete theory of consciousness.

A complete understanding of consciousness will not be

possible until we have a neurophysiological explanation of the behaviors characteristic of conscious beings, plus a metaphysical account of mind which tells us what conscious experiences really are. Nowadays it is easy to forget that if we had the scientific theory, we would not automatically have the metaphysical account, which perhaps is more important for the conduct of human life. But even if we had such a pair of theories, this would not change facts that things *seem* this way or that. For my present purposes, the latter facts are all that matter.

An Aside for Philosophers of Mind

For philosophers of mind, a certain technical objection may have suggested itself as early as the beginning of this chapter. This objection arises from arguments which suggest that seeming is reducible to some other psychological or physical phenomenon, or at least has borderline cases with such a phenomenon. A prime example of such an argument is Dennett's attempt to explain the ways things seem to people in terms of "discriminations" in the brain, which Dennett likens to "'judgments'."¹⁹ According to Dennett's view, there are no "seemings" separate from these.²⁰ Such an hypothesis might make one wonder whether there really is a kind of "seeming" which is just a matter of being conscious of things, apart from belief or judgment about one's surroundings. But even if Dennett's view of seeming as judgment were right, it would be

irrelevant to what I have said about the meaning of "to seem." There still would be a difference between judgments of the kind which we ordinarily regard as seemings, and the very different (less compelling?) judgments which we usually regard as acts of belief. For example, many people believe, fervently and unshakably, that there is a God, without *experiencing* that there is a God as a mystic might claim to do. A blindfolded person can believe that an apple is red without it *seeming* to him or her that that apple is red. Many people believe unshakably that the Earth is round, without the Earth ever *seeming* round to them. Views like Dennett's also may imply that there are borderline cases between seeming and belief or judgment.²¹ But such borderline cases pose no threat to anything I have said, since there still would be some clear-cut examples of seeming.

No fact about the neurophysiological basis of consciousness can change the fact that a conscious being's surroundings seem, or appear, some way or other to that being. Knowledge about the nature of consciousness may change our understanding of that fact, and perhaps even change our psychological feel for it²² — but those are different kinds of change.

Subjective Fact and Intentionality

States of mind in which something seems some way or other involve the apparent truth of facts and/or the apparent existence of objects. Hence such states are states of a

special sort; they are what philosophers of mind call *intentional* states. An intentional state, according to one widely used characterization, is "about" something; it has an object of some sort. Commonly invoked examples of intentional states include states of sensation, in which some object is sensed, and states of thinking, in which some subject matter is being thought about. A related notion, also widely used in the philosophy of mind, is that of *content*. I do not wish to discuss the relationship between intentionality and content in detail; for our purposes, it is enough to say that an intentional state is a state with content.²³

Philosophers who think about intentionality typically regard it as one of the most central features of consciousness. Earlier I characterized consciousness in terms of subjective facts without mentioning intentionality. However, my characterization directly implies that consciousness is intentional. Conscious states have apparent facts and entities as intentional objects.

There may be more to an intentional state than its associated subjective state. One can, for example, continue to believe a fact while not currently thinking about that fact.²⁴ Nevertheless, the instances of seeming associated with conscious mental states have subjective facts and beings as intentional objects, and therefore are intentional and insure that the conscious states with which they are associated are intentional.

In what follows, I sometimes will speak of the "content" of an instance of seeming. By that I will mean the set of

subjective facts and beings associated with that instance.

The Subjective and the Objective

Earlier I pointed out that the notions of subjective fact and subjective being suggest a certain parallel between seeming and truth. Now I will explore this parallel more thoroughly, and will touch upon some deeper issues connected with it.

One can think of facts which seem to be the case as having *subjective truth*, in contradistinction from truth as such, which is *objective truth* (independent of viewpoint). For example, if a fact is the case for one instance of seeming but not for another instance of seeming, then one can think of that fact as being subjectively true for one instance of seeming but not for the other. Similarly, one can contrast subjective being with being as such, which is objective being. Thus, an entity which exists for Henry but does not exist for John can be thought of as existent with respect to Henry and as nonexistent with respect to John. One can think of subjective fact and subjective being as *subjective* or *perspective-dependent* notions of being and of truth, in contradistinction from the familiar *objective* or *absolute* notions. This way of thinking about subjective fact really adds nothing to the notions of subjective fact and being, but it will prove suggestive later on.

Using the terminology of subjective truth and being, we can say that consciousness is a phenomenon in which

subjective truth plays an essential part. There is consciousness if and only if there are facts that are true *for* something, not merely objectively true. The presence of consciousness in the world is the presence of subjective truth in the world.

Of course, the idiom of subjective fact and subjective being is at bottom simply a way of talking about phenomena involving seeming. It cannot conflict with other ways of talking about phenomena of this sort.

Relativism and Subjectivity

Note well that the use of the notion of subjective truth does not back us into a position of relativism with regard to truth. I have deliberately avoided using the term "relative" to describe this notion since I do not want my intentions to be misunderstood. The idea of subjective fact does not replace objective truth with something relative or subjective; on the conceptual level, there is enough room in the universe for both subjective truth and objective truth.

The fact that there is subjective truth does not, by itself, imply that there is any truth beyond the subjective. If one believes in relativism, then one consistently can deny the reality of objective truth while still admitting the reality of subjective truth as a special kind of relative truth. However, one does not have to believe in relativism to recognize that subjective truth is real. *And not even a committed relativist can deny that there are ways things seem.*

Every argument which relativists might offer against the objective reality of the world is, in effect, an argument that there *only seems* to be an objective world. If there did not seem to be an objective world, then relativism with regard to truth would not be controversial: no one would doubt that truth is relative to perspective. An argument intended to persuade someone to adopt relativism must represent an attempt to convince an audience to ignore how things seem to them — to ignore the appearance that there is an objective world. Thus, relativism, through its rhetoric and also through its very existence, acknowledges the reality of seeming. This acknowledgment is its first act.

It is worth noting that subjective fact and subjective being are, in a deeper sense, *objective* phenomena. If it seems in a particular instance that there is a loud noise, then it *really is* the case that it seems in that instance that there is a loud noise. Thus relativism, by tacitly presupposing the existence of ways things seem, unwittingly bases its claims upon a foundation of objectivity.

In later chapters (4 and 13) I will address issues of objectivity and relativism in more detail, and will explore further the relationship between subjective and objective truth.

Note also that the notion of subjective fact does not have built into it any prejudices about what subjective facts really are. A postmodernist might want to claim that subjective facts are of a merely linguistic or "textual" character. I have not ruled out this possibility, nor have I endorsed it. A philosopher who claims that the content of our experience is

built on social factors, belief systems, theories, values, or the like will not be able to use these claims to attack the notion of subjective fact — for the notion of subjective fact presupposes nothing about the real nature or origin of such facts.²⁵ A subjective fact is simply a fact about how things seem — regardless of what a fact about how things seem ultimately turns out to be, and regardless of how it originates.

Subjective Fact as a Fundamental Notion

Throughout the rest of this book, I will treat the notion of subjective fact as a fundamental notion. I will use it as a basic concept for explaining other concepts, and will define other concepts in terms of it. This use of the notion of subjective fact will encompass the use of the notion of instance of seeming, which is a component of the idea of subjective fact. I will use these twin concepts in various examples and arguments about how things seem. With the help of these examples and arguments, I will arrive at some conclusions about the objective world, thereby accomplishing the project of this book and doing some philosophy along the way.

Taking subjective fact as a fundamental notion amounts to taking the notion of *seeming to be the case* as a fundamental notion.

This decision to adopt subjective fact as a fundamental notion is not meant to suggest that I think the concept of

subjective fact is undefinable. Already I have defined this concept, albeit non-rigorously, by presenting an informal definition and several examples. My decision to treat this concept as fundamental also does not imply that I think it is unanalyzable. Anyone is free to argue that consciousness — and its central feature, subjective fact — really is identical to some neurophysiological or behavioral phenomenon, or to something that happens to Cartesian egos, or whatever. Anyone is free to maintain that some other notion is more fundamental and to define subjective fact in terms of it. By adopting subjective fact as a fundamental notion, I am simply choosing a conceptual starting point for further argumentation. This particular starting point is extremely convenient, since it allows us to make statements about how things seem without importing any major assumptions about what really exists.

Chapter 3

The Happenings Within

In the last chapter, I developed some conceptual machinery for the study of consciousness, or at least of consciousness' central feature, the way things seem. In this chapter, I will begin to explore the relationship between consciousness and time. I will develop further the concept of an instance of seeming, and will point out reasons to believe that instances of seeming are *events* of some sort. Since events happen in time, this exploration will begin to shed light on the subject of time.

Later in the book, I will draw several other conclusions about time — some of them perhaps quite surprising. For now, I will concentrate on the apparent temporal features of instances of seeming.

Consciousness Events

Instances of seeming are rather puzzling items. On the face of it, they do not appear to belong to any of the standard

categories of entities recognized by philosophers, such as the categories of things, events, and abstract objects. It could turn out that instances of seeming form an entirely new metaphysical category. However, there is some reason to think that instances of seeming are *events* of a sort.

Consider how things seem to you now. When you consider this, you find out at once that things seem a certain way. If you continue to register how things seem to you, you also find out that things are continuing to seem that way, or else are seeming other, new ways. Something is *happening*; things *are seeming* some way or other. Apparently, things' seeming some way or other is not only a static fact; it is something that happens. Hence, *prima facie*, when things seem a certain way an event is happening — an event of appearing, of things seeming some particular way. An event of this sort is a single example or occasion of things seeming a certain way. Hence it is plausible to identify such an event with an instance of seeming. Consequently, it is plausible to suppose that an instance of seeming is an event.

The preceding highly informal argument is far from conclusive, but it is suggestive. It does not pretend to establish once and for all that an instance of seeming is an event. Also, this argument does not tell us what kind of event an instance of seeming might be. A decision upon this matter would have to depend upon which philosophical view about the nature of mind is correct. If a materialistic or physicalistic view is right, then all mental events are physical in nature; hence an instance of seeming, if it is an

event, can only be a physical event. Presumably, this physical event is the very complex event or process which your brain undergoes, by virtue of which you are conscious of something. Different physicalists will have different ideas about what that event is. If the dualistic view of mind is right, then an instance of seeming is an event that happens to a non-material soul or self; it must be whatever happens to your soul, by virtue of whose happening you experience something. If behaviorism is right, then there really is nothing to mental life besides observable behavior. In that case, an instance of seeming might be something like the event of the onset of certain bodily behaviors. However, we do not need to ascertain the real nature of this event to continue with our project.

An event in the history of a dead lump, or of a zombie (an imaginary being with no subjective life),¹ does not involve things seeming any way. Some events in the lives of humans (in particular, some mental events) are associated with things seeming certain ways. Hence some events in the lives of humans are events of the sort which it is tempting to identify with consciousness events. Events in the careers of lumps and of zombies definitely are not consciousness events.

The above arguments suggest that we may think of an instance of seeming as an event of a particular kind — an event whose occurrence consists in things seeming a certain way. However, the concepts I have introduced for talking about instances of seeming do not force us to think this way. For now, we can regard an instance of seeming simply as an

inherent feature of subjective facts. Later, we may begin to regard all instances of seeming as events, if we have sufficiently good reason to do so. If we do that, then we can treat "In the instance __, it seems that __" as an open sentence in which an event name and a sentence must be inserted.

We have not shown conclusively that instances of seeming are events. Be that as it may, I am going to use the term *consciousness event* as a synonym for *instance of seeming* throughout this book. This term is intended to highlight the appearance of happening which accompanies instances of seeming. If consciousness events turn out not to be events, then this terminology will be misleading, but no more so than some commonplace expressions like "plastic flower." Later on, the term "consciousness event" will turn out to be quite fitting for several reasons.

Consciousness Events: Some Examples²

A consciousness event is an instance of things' seeming to be some way or other. Right now, your surroundings seem a certain way to you. A very short time later (when you read the next word, hear a new sound, or change your view slightly), your surroundings seem a slightly different way. Each time things seem to be some new way to you, there is another consciousness event.

The following examples of consciousness events are meant to show how consciousness events enter into our everyday experience. They should enable the reader to get a

better psychological feel for the concept of a consciousness event. Also, they point up the apparently "dynamic," or event-like, nature of consciousness events.

(1) Look at the period at the end of this sentence, then suddenly look at something else. When you move your eyes, you pass from one consciousness event to another. One instance of seeming has ended in time (at least in time as you feel it); another instance of seeming has begun. What seems to you to be the case after the change is not what seemed to you to be the case before the change.

(2) Visualize a yellow square. Then stop visualizing it; just visualize blackness instead. When the square disappears, you pass (at least in apparent time) from one consciousness event into another. Your subjective "world" is not the same after the change as before.

(3) Look at the period at the end of this sentence; then keep on looking at it. After the first brief moment of looking, your experience of the period is not quite the same as it was during that first look. Among the differences: You now *remember* having looked at the period, and your thoughts, moods, and other sensations (such as the sensation of the background noise in the room) may have changed slightly. You have passed into another consciousness event.

From Brain to Cosmos

(4) As you read the word "as" which began this sentence, you were having one consciousness event. *Now* you are having another (and perhaps you had several more since you read the "as"). (Pause here.) And now you are having still another. (Pause.) And another. (Pause.) And...

You can find other examples for yourself. It is not difficult to identify consciousness events in your own experience.

Note that the finding of consciousness events for yourself does not require you to engage in any genuine *introspection*, or observation of your own inner processes. You only need to describe, or take note of, how things seem; observations of apparent structures or processes within yourself are quite unnecessary for this. Finding a consciousness event is simple; ascertaining that there is more than one consciousness event requires a slightly more complicated act. To do the latter act, describe to yourself how things seem, or take note of how things seem, and then do the same thing again a moment later. If what you noted the second time is different, then you can infer that there were two different consciousness events. (Of course, this last inference may require you to rely on memory of the first of the two descriptions, and hence may depend upon the tacit assumption that your memory of the previous consciousness event was correct. In Chapter 4 we will see that this assumption is not always necessary, especially in cases like those of examples (1)-(4) above.) Finding consciousness

events is easy, once you know what to look for. It is not a matter of spotting some rare bird in the forest.

Consciousness Events and Experience

You, the reader, are conscious now. The present condition of your consciousness involves all that *seems* to be happening to you right now — what you are sensing, thinking, feeling, and so forth. This present occasion of your being aware of your outer and inner surroundings is an instance of seeming. Speaking loosely and intuitively, we can say that the fact that this instance of seeming exists is the fact that there exists, for a brief time, a particular "realm" of subjective facts. These subjective facts are associated with thinking, sensing, and other conscious processes. After a fraction of a second, you will be undergoing a different consciousness event. You still will be conscious, but this time you may be conscious of different things.

Ordinarily, a consciousness event is not merely a single instance of one kind of seeming — the kind we associate with seeing, hearing, thinking, or the like. If I am feeling a pain, then my present awareness includes other things besides that pain — for example, I may be thinking about philosophy as well. If so, then I am undergoing a consciousness event that involves both of these kinds of seeming — the ones associated with the mental phenomena which we ordinarily call "feeling pain" and "thinking." The subjective facts which are the case in a single instance of

seeming may be of the sort which arise from sensing, thinking, feeling, or any any of the other phenomena which can help to make up a single instance, or incident, of being conscious.

Consciousness events are not the same as experiences. You can have an experience which involves the content of several of your consciousness events — for example, the experience of reading a sentence in a book. Also, you can have several different kinds of experience, and hence (arguably) several different simultaneous experiences, during the same consciousness event — as in the preceding example of feeling pain while thinking. From now on I will use the word "experience" to refer to those conscious mental events which we customarily call experiences. If an experience involves the content of a consciousness event, I will speak of that experience as "happening during" the consciousness event. I will not address the question of whether experiences might be entirely reducible to consciousness events, or the related question of whether, in some special cases, consciousness events might count as experiences. The answers to these questions are not necessary for my project.

The Timing of Consciousness Events

In ordinary human life, a consciousness event appears to be a brief event in the history of a conscious subject. This fact suggests another way to point out what a consciousness

event is. A consciousness event might be thought of as a *specious moment* in the conscious life of a subject — that is, a subjective moment of experience, roughly the same as what C.D. Broad termed "the Specious Present."³ It is intuitively reasonable to characterize a consciousness event as a specious moment of conscious life, while recalling that this moment need not be instantaneous in clock time.

This last point bears repeating: it is important to remember that consciousness events need not be instantaneous. Empirical evidence points to the conclusion that a consciousness event takes a substantial amount of time — typically about a third of a second. Several experiments strongly suggest that a stimulus received within about 300 milliseconds of an earlier stimulus can affect the conscious perception of that earlier stimulus.⁴ This suggests that the conscious perception of a stimulus takes at least this long to occur.⁵ (Even if a single consciousness event lasts this long, some consciousness events might *seem* to be quicker.)

Unusual Consciousness Events: Some Possibilities

In ordinary human experience, consciousness events appear to be event-like "specious moments" of the kind described above. However, the notion of a consciousness event does not fix most characteristics of a consciousness event. To underscore this fact — and also for other reasons

which will surface later — I will mention some possible kinds of consciousness events quite different from the ones humans usually undergo. I am not claiming that consciousness events of these kinds really are possible, physically or otherwise; perhaps they are possible, perhaps they are not. My only point in describing these "possibilities" is to show that the notion of a consciousness event does not exclude consciousness events quite different from those which people normally have.

One example of an unusual consciousness event would be an isolated consciousness event which is not part of the life of a persisting subject. This would be reminiscent of, though not identical to, the stray perceptions — perceptions that belong to no one in particular — which Hume once considered.⁶ Another possibility is a consciousness event which is not transitory — which is not an event that *happens* in the usual way, but which simply *is*. This possibility becomes less impossible if we recall that a so-called consciousness event is not defined to be an event. (Theologians interested in the concept of eternity have suggested that God exists beyond the transitory sort of time that we know.⁷ We should leave open the possibility of non-transitory consciousness events, if only because we do not want to beg any *theological* questions.) Still another example would be a consciousness event which occurs in the absence of physical events like those in the brain. We have not ruled out the possibility of such a consciousness event, nor have we ruled out the opposite view that all consciousness is physical.

We have not ruled out the possibility of consciousness events whose *timing* is very different from that of ordinary consciousness events. For example, we have not asked whether it is possible for there to be an instantaneous consciousness event. (Those of humans, as I have said, are of finite duration.) Nor have we asked how the consciousness events which happen to a single subject are arranged in time. In subsequent chapters I will address some questions much like these, about the possible kinds and relationships of consciousness events.

It is not conceptually necessary that a consciousness event involves any of the familiar features of consciousness, except those which follow from the mere existence of a way things seem.

The Logic of Consciousness Events: Three Appendices

The following three appendices deal with some logical properties of subjective fact. They use standard ideas of modal logic.⁸ The reader who is not a logician can skip the details, but should be aware of two key conclusions, which I will describe here before beginning the appendices. The conclusions are that consciousness events possess logical features which we may call *logical incompleteness* and *worldlike character*. (The term "logical incompleteness" is borrowed from mathematical logic, where it denotes a property of formal systems somewhat analogous to the

property of consciousness events which I will discuss.)

By saying that consciousness events are logically incomplete, I mean that it is possible for a sentence (or proposition) and its negation both to fail to be the case for a consciousness event. It is quite possible for P and not-P both to fail to be the case for a consciousness event. For example, if I am not looking up into the sky, then it is not the case for my present awareness that Pluto is overhead, but it also is not the case for my present awareness that Pluto is not overhead. Pluto's present position simply is not a part of my inner world right now. (Recall the Pluto example from Chapter 2.) This feature of consciousness events may seem rather trivial, but it will prove quite useful in later chapters, where I will use it in an investigation of the problems of personal identity, the unity of the self, and the nature of the unconscious mind.

By saying that consciousness events have worldlike character, I mean that a consciousness event is analogous, in some important respects, to what philosophers call a *possible world*. If a sentence (or proposition) P is the case for a consciousness event, then one can think of P as being true *at* that consciousness event, just as a modal logician might say that P is true at a possible world. From a logician's point of view, a consciousness event is much like the possible worlds, or alternative possible universes, used in metaphysics and in modal logic. (Some points of similarity and of difference are mentioned in the appendix.) This analogy between consciousness events and worlds does not stretch the truth too far, since a consciousness event is

associated with a subjective "world" of subjective facts and beings.

An intuitive recognition of the worldlike character of the subjective realm may well have lain behind Leibniz' analogy between the monad (or perceiving entity) and the cosmos.⁹ My view of consciousness events as worldlike is reminiscent of this Leibnizian view. (Of course, my view does not imply most of Leibniz's other ideas about monads, and is based upon simple logical properties of consciousness events rather than upon strong metaphysical arguments.)

Appendix A: Subjective Fact and Modality

From a logician's standpoint, the notion of being-the-case-for is a *modality*, with its own special logical characteristics.¹⁰

The locution "For x, it is the case that" introduces an intensional context in sentences of the form "For x, it is the case that P." The truth value of such a sentence is not determined by the truth value of P. One can regard "For x, it is the case that" as a modal operator. Because this operator requires an object (a value of x) as well as a formula to act upon, this operator resembles the modalities expressed by "knows that" and "believes that" more closely than it resembles the modalities of necessity and possibility discussed in elementary modal logic texts. We might call the operator "for ____ it is the case that" the *being-the-case-for operator*, or the *BTCF operator* for short. The modality

which it represents — that of subjective truth — can be thought of as a *subjective alethic* modality.

One could formalize the notion of subjective fact and develop a system of modal logic for the BTCF operator. I will not undertake such a formalization here because it will not be necessary for what follows. However, I will make a few remarks on the logic of the BTCF operator.

The syntax of the BTCF operator is clear from the preceding remarks. The semantics of the BTCF operator follow from my previous discussion of subjective fact. One can state the truth condition for the wff "For x, it is the case that P" as follows: "For x, it is the case that P" is true if and only if "x" is a tag for an instance of seeming in which it seems that P. This truth condition makes the truth of "For x, it is the case that P" depend solely upon facts about how things seem — specifically, upon what seems to be the case in what instance, and upon nothing else.

Appendix B: Consciousness Events as Incomplete Worlds

One can think of a consciousness event as that which fills in the blank in the modal operator "for ____ it is the case that." Alternatively, one can think of a consciousness event as the analogue, for subjective fact, of the important philosophical concept of a *possible world*. One can regard "For x, P" as asserting that P is true *at* the consciousness event x, in much the same way that a sentence or proposition

can be said to be true at a possible world. This is the basis for my earlier assertion that consciousness events have worldlike character.

Consciousness events, conceived of as analogs of worlds, make sentences or propositions true or false. One could construct a modal semantics in which consciousness events play the part of worlds. Consciousness events differ from worlds in respect of their logical incompleteness; a consciousness event need not make every statement about experience either true or false. Consciousness events behave like what philosophers have called "incomplete worlds" or "possibilities",¹¹ rather than like standard possible worlds.

Truth at a consciousness event creates an intensional context. The intensionality of the BTCF operator (discussed in Appendix A above) reflects this. The familiar intensionality of mental contexts¹² arises at least in part from this intensionality — the ability, as it were, of consciousness to endorse propositions. For example, "John believes that P" tells us nothing, in the absence of other information, about whether P is true. But it does tell us something about John's subjective world. The intensionality of consciousness events lies behind at least some instances of false belief (consider beliefs arising from hallucinations and illusions). If "John knows that A" fails to follow from A, this is largely accountable to the character of John's experience. (If John were aware of everything — that is, if subjective truth for John's consciousness events coincided with objective truth — then this occurrence of "knows" might not introduce an intensional context.) The

intensionality of mental contexts arises, at least in part, from the conceptual connection between mental acts and subjective fact. Hence this intensionality arises, at least in part, from the very essence of consciousness.

Appendix C: Are Consciousness Events Logically Consistent?

It is tempting to suppose that consciousness events, like possible worlds, are *logically consistent*. This will be the case if for any value of "x", the sentences "For x, P" and "For x, not-P" are not both true. The claim that this is the case is intuitively appealing, and is correct for our ordinary experiences. It also is plausible in view of our ideas about consciousness events. We think of a consciousness event as being associated with a realm of subjective fact, and particularly with a *unique* realm of subjective fact. According to this picture, any given subjective fact either belongs to the realm (making "For x, P" true), or does not belong to the realm (making "For x, not-P" true), but not both. If this is a conceptual truth about consciousness events, then consciousness events are logically consistent.

In this book, I will not address the general question of the logical consistency of consciousness events.

Chapter 4

On Knowing What Just Happened

In the last two chapters I introduced some concepts and devices for the description of experience. My first application of these devices will be to a problem about knowledge. This is the problem of justifying beliefs about the past. With the help of the concepts of subjective fact and of consciousness events, I will argue that some of one's knowledge of the past can be known with certainty. More precisely, I will contend that there are certain facts about immediate past experience which are true if and only if they seem true. Such facts can (with certain provisos) be known with the same degree of certainty as facts about the way things seem now.

Later in the book I will discuss other questions about the trustworthiness of experience (such as experience of other minds and of physical objects). These problems are beyond the scope of this chapter.

Certain Knowledge About Experience

In most cases, the fact that something seems to be the case does not guarantee that it really is the case. We may ask whether there are any statements whose truth is implied by their apparent truth: that is, statements which, if they seem to be the case, are the case. The existence of statements of this sort would imply that some of our knowledge is, at least in a restricted sense, certain — not merely highly probable, as is most of our knowledge. The main objective of this chapter is to show that there are statements of this kind.

When attempting to show that there are statements of this sort, we are better off not using any premises which are less certain than the knowledge whose certainty we seek to evaluate. If we use such premises, the argument for the certainty of the knowledge in question becomes uncertain, undermining our sought-after confidence in the certainty of the knowledge. Although an argument whose premises are truly certain may be an unattainable ideal, we will aim toward this ideal by building the argument on facts about what seems to be the case.

In addition to restricting our premises in this way, we should restrict our forms of argument in certain ways. For example, we should not employ inductive reasoning, or any other technique of argument or proof which, by its very nature, can lead from more certain premises to less certain conclusions. We also may exclude the use of certain

deductive arguments, as I will point out shortly. If we are thorough, we even will restrict the sizes of the arguments we use: those arguments cannot be too long!

One's confidence in the validity of a very long argument depends upon one's beliefs about the reliability of memory. If an argument is too long to keep in mind all at once, then one believes its conclusion because when one reaches the conclusion, one recalls that one went through the earlier steps and found them correct. (This kind of recall may be aided by the making of notes — either actual notes on paper or so-called "mental notes"). When one finally gets to the conclusion, it seems to be the case that the argument had earlier steps which were correct. But one does not now remember those steps all at once; if the argument is long enough, one cannot hold all the premises in memory now, together with enough steps to recognize the argument as formally valid now. Hence, if we aim for the highest possible degree of certainty, we should not assume uncritically that a very long argument really is correct. (Even if you are not a Cartesian skeptic, certainly you must misremember things at times.)

Short arguments sometimes can escape this difficulty. When thinking about very short arguments, one can (and often does) hold the essential features of the whole argument in memory all at once. Metaphorically speaking, one "sees" that the argument is correct. Thus, if one wants to prove to oneself that a particular kind of experience is reliable, one must use short arguments of this kind to do so. Any argument which claims to show that certain sorts of

experience are completely trustworthy must be so short that we can grasp the argument all at once *now*, in a single mental act. (If the understanding of the argument requires some preliminary thinking which is not part of the argument, then that is acceptable too.)

Some general claims about subjective fact can be grasped *now* in just this fashion. Consider the claim (made in Chapter 2) that "Pluto is overhead" is neither the case for you now nor not the case for you now. It may take some time to understand the concepts necessary to grasp the meaning of this claim. It may even take some time to convince yourself of the reasons given in Chapter 2 for accepting this claim, or to find your own reasons for accepting it — if you are going to accept it at all. But once you understand the claim and become sufficiently familiar with the argument for it, you can grasp this claim, together with its main supporting reason, *all at once right now*. One who has read and understood previous chapters, and who understands the Pluto claim and the arguments for it, could explain and justify this claim without having to rifle through the earlier pages of the book. Once one grasps the concept of being the case for, the actual argument necessary to establish the claim in question is not that long. It can be summarized in one informally worded sentence: "Pluto doesn't seem like *anything* to you; it just *isn't there* for you." One can grasp this kind of justification *now*, without having to cite past arguments whose many steps one no longer remembers.

There is nothing wrong with using long arguments to lead

up to a claim whose justification can afterwards be grasped without the aid of such arguments. Without such longer arguments we would get nowhere. But since we are trying to derive our conclusions from facts about how things seem, we must try to use only arguments which can be grasped "all at once" — as when one "sees" a brief argument after following it through and understanding all of its steps.

Seeming to Seem

It often *seems* to an observer that something *seems* to be the case. The following examples illustrate this fact.

(1) I am looking at an orange square. Someone asks me "What seems to be in front of you?", so I begin to pay attention to what I am seeing. I realize that it seems to me that there is an orange square. It seems to me, at the moment of this realization, that it seems to me that there is an orange square.

(2) Does there seem to be a book here? It seems to you, at the moment when you realize that you are reading a book, that it seems to you that there is a book in front of you.

(3) Look straight ahead. Then suddenly look away at some other object. The moment after you looked away, it seems to you that something has just

happened. Statically looking at the new object is not the same experience as looking at the new object after just having moved your eyes. In the latter case, you get the impression that something has just taken place. This feeling is not present if you merely stare at the new object.

(4) You hear a sudden, loud noise resembling a gunshot. At once you are aware that "something happened." You didn't just hear the noise — you were "hit" by the noise sensation, then became aware that something happened. Before you had time to think "Hey, what was that?", you became aware that something had occurred, that things had altered. What happened was that you perceived a noise. (The external physical event that caused the noise also happened, but it happened before you heard the noise; the experience resulted from a later perceptual process.) Although you did not yet have time to think about it, it seemed to you that you just had an experience happen. In other words, it seemed to you that something seemed to be the case.

I will discuss examples (3) and (4) first, and will return to (1) and (2) later.

In examples (3) and (4), an experience seems to have occurred just a moment ago. In (3) the experience was one of a change within the visual field; in (4), it was one of a loud auditory bang. These examples can be redescribed in

terms of subjective fact. When they are redescribed, they come out as examples in which something seems to have just seemed, in the immediate past, to be the case. In (3), it seems to be the case for you now that a different view seemed to exist. In (4), it seems to be the case for you now that there seemed to be a loud noise.

In both of these cases, it seems *to you now* that some fact was, or is, the case for you — but not that that fact is the case for you *right now*. It seems to you now that that fact was the case *for you sometime* — specifically, in the immediate past. In (3) or in (4), the relevant fact (the occurrence of the other view or of the bang) does not seem to be the case for you *now*. It seems to have been the case for you in the *immediate past* — in a preceding moment of time.

We can restate these examples in the language of consciousness events. In example (3), it is the case for your present consciousness event that there is a consciousness event for which certain facts are the case. (In the preceding sentence, I used "there is" to cover, not only putatively present consciousness events, but putatively past and future ones as well. I can indulge in this practice because the quantifier over consciousness events has been interpreted substitutionally (recall Chapter 2); hence quantifying over past and future consciousness events does not commit us to the reality of past or future entities. The reader who is uncomfortable with this usage of "there exists" may replace this "there exists" by "there was, is, or will be.") In example (4), it is the case for your present consciousness event that

there is an instance of seeming (consciousness event) in which there is a loud bang.

These examples illustrate an interesting fact about experience: that it can seem to be the case that an instance of seeming *just happened*.

Experiences of this sort must not be confused with experiences involving memory. It may seem to me at 1 pm that I went to work this morning. But in that case, it does not seem to me that my arrival at work just happened. At 1 pm, none of the consciousness events which make up my experience of arriving at work seem to me to exist. They are no longer part of my inner world — as the older visual perspective in example (3), or the loud bang in example (4), still are for a moment after they happen. Only the memory of what happened this morning remains; I can no longer notice myself experiencing my arrival at work, as I can now notice myself experiencing the loud noise which I heard a fraction of a second ago.

The consciousness events which I had while arriving at work much earlier in the day do not exist for me at 1 pm. But in example (3), the consciousness event for which the previous view exists is part of the subjective realm of the next consciousness event. The fact that something was just experienced (that is, that there was a consciousness event for which something was present) is the case for the next consciousness event.

Examples (1) and (2) are much like examples (3) and (4). In them, it seems to a subject that in some instance of seeming, things seem a certain way. (Of course, the subject

might never think of describing his or her experience in this way, but the relevant instance of seeming exists for him or her nevertheless.) In examples (1) and (2), it is not so clear whether the consciousness event which seems to exist seems to be *immediately past* or seems to be *present but being replaced*. But in either case, the apparent transitoriness of the instance of seeming is evident. The general sort of experience involved is the same as in examples (3) and (4): it seems to me now, that something or other seems (or seemed) to me to be the case.

The Redundancy of Seeming

One of the most notable *logical* characteristics of seeming is this: if it seems to be that things seem a certain way, then things do in fact seem that way. (In symbols: if it seems that it seems that P, then it seems that P.) This putative fact has been discussed in the literature, and has been made quite explicit by Dennett.¹ Here I argue that this putative fact really is the case, and I also explore the reasons why it is the case. The following example illustrates this characteristic of seeming.

Suppose that I am now seeing an orange square. A Cartesian skeptic, or a behaviorist, comes along and says: "I admit that you are having the experience colloquially called 'seeing an orange square.' You claim that it seems to you that there is an orange square, and I concede that you are neither psychologically nor linguistically confused nor lying.

Yet it does not really seem to you that there is an orange square. Instead, it merely *seems to you* that it seems to you that there is an orange square."²

An appropriate informal reply to this skeptic runs as follows. When I say that it *seems* to me that there is an orange square, I am, of course, describing the way things appear to me. Consider what the skeptic is saying when he says that it *only seems* to me that it seems to me that there is an orange square. The positive part of what he is claiming amounts to this: that it appears to me as if an orange square were apparent to me. But how can this be true unless an apparent orange square does in fact figure in my experience in some way? Even if I am easy to deceive, I cannot be fooled in such a way that it seems to me that there seems to be an orange square, unless an apparent orange square plays some part in the experience which is deceiving me. To manipulate my experience so that it seems to me that I am experiencing an orange square, you would have to introduce an orange square into my experience in some way or other. If you did not do this, then it could not even seem to me that there *seems* to be an orange square. But if you did do this, then it really *would* seem to me that there is an orange square. If there is no orange square at all in my experience, then how could it seem to me that there even *seems* to be such a thing? Evidently, the mere *appearance* that there *appears to be* an orange square involves the *appearance of an orange square*.

One may formulate this argument slightly more rigorously as follows. Let P be some sentence; for the first

stage of the argument, let P be a sentence asserting the existence of an object of a particular sort, say the Eiffel Tower. This sentence expresses a singular existence claim; as such, it is fraught with philosophical problems. However, these problems do not affect us here; the important point is that it may seem that P is true, as it would, for example, to a tourist looking at the Eiffel Tower. Suppose that it seems to be the case that it seems to be the case that P. Rewriting this in terms of instances of seeming, we get: in some instance of seeming x, it seems, in some instance of seeming y, that P. But this implies that in instance x, the instance y seems to exist. (Otherwise it could not seem in x that anything was true of y.) Now suppose that P is not the case for x. Then for x, there is no Eiffel tower. There simply is no such object for x; the "inner world" or "point of view" embodied in x just does not contain the Eiffel Tower. Thus, for x, it cannot seem to be the case that there is a consciousness event for which such a tower exists. No fact about the Eiffel Tower can seem to be the case in x, for in the point of view embodied in x, no trace of the Eiffel Tower is to be found. The consciousness event y may exist for x, but it cannot be the case that for x, there is a consciousness event for which the Eiffel Tower exists. For x, y cannot be such a consciousness event.

For the second phase of the argument, let P be an arbitrary sentence which may be true or false, but which may seem to someone to be true. For P to seem true to you, it must seem to you that a certain situation obtains; for example, if P is "It is hot in here," then it seems to you that P

if and only if it seems to you that your surroundings are hot. (If the apparent truth of P is not associated with an apparent situation in this way, then P could not seem to be the case to anyone; there would be nothing that it is like to experience that P.) Suppose that it seems to be the case that it seems to be the case that P. Once again, we rewrite this using two consciousness events: in some instance of seeming x, it seems that in some instance of seeming y, it seems that P. Now suppose that P is not the case for x. Then for x, the situation associated with P does not obtain; for x, there is no such situation at all. It follows that it cannot seem in x that there is a consciousness event for which the situation associated with P obtains. There simply is no such situation for x. Therefore, if P is not the case for x, then for x, there is no consciousness event y such that P is the case for y. (There may be a consciousness event y which exists for x, and is such that *in fact* P is the case for y. But this would not imply that for x, P is the case for y.)

This argument points out an important feature of the phenomenon we call seeming: namely, that what seems to seem to be the case, seems to be the case. The argument rules out the possibility that something "only seems" to someone to seem to be the case. Of course, the argument does not rule out mistakes about how things seem. It allows that many mistakes of this sort are possible, or arguably are possible. For example, I might be fooled into believing that it has seemed to me that there is an orange square, when in fact it has not so seemed. (Things of this sort happen in typical cases of erroneous memory.) However, mistakes of

this kind are irrelevant to the argument of the preceding paragraphs. The fact that it seems to me that there seems to be an orange square presupposes the fact that there *does* seem to be an orange square.

The arguments of the last few paragraphs issue in one conclusion: if it seems to you that it seems that P, then it really does seem that P. The claim that the appearance of something is merely an appearance of an appearance, as opposed to a *real appearance*, betrays conceptual confusion. There can be no difference between an "apparent appearance" and a "real appearance," since all appearances only consist in what is "apparent." (Recall endnote 1.)

If things "really seemed" one way but appeared to seem some other way, then that *other way* would be the way things really seem to you. The other way would be the appearance that you actually get. If the fact P only seemed to seem to be the case, then P would play a part in the way things seem — in how things seem to you at the moment. Hence P would in fact seem to be the case.

If we recall the discussion of the consciousness sense of "seem" in Chapter 2, we can see where the above conclusion really comes from. It is a partial rendering of what we mean by "to seem" in the consciousness sense. We can best understand this by thinking of seeming as a modal operator, as we did in the appendixes to Chapter 3. The following paragraph (which non-logicians may skip over) shows this.

Let S be a "seeming operator"³ — that is, an operator such that for any sentence P, SP is true if and only if it seems that P (either simply, or in some particular instance).

Let M be some unspecified modal operator. Suppose that SMP does not entail SP. Then it can seem that MP even if P does not seem to be the case. If MP can seem to be the case whether or not the fact that P plays any role in experience, then "MP," by itself, cannot convey information about the role played in anyone's experience by the fact that P. Such an operator M cannot adequately represent seeming or appearance. A modal operator of this kind may be interesting and important — the past tense operator of tense logic is one obvious example⁴ — but it cannot be an operator representing seeming. To qualify as a seeming operator, a modal operator must not give an apparently true result when applied to a sentence which has *nothing at all* to do with what is being experienced.

This argument shows that the rule that "what seems to seem to be the case, really seems to be the case" is a general logical property of consciousness. It is one of the criteria which a sense of "to seem" must meet to be considered identical to the consciousness sense of "to seem."

We can put this rule in the language of subjective fact (and avoid further repetitions of "seem" and "appear") as follows. Let P be a sentence. If for some consciousness event, it is the case that it seems to be the case that P, then it seems to be the case that P. Proceeding further, we can replace "it seems to be the case that" with the strictly equivalent "it is the case for some consciousness event that." Then we get: If for some consciousness event it is the case for some consciousness event that P, then it is the case for some consciousness event that P.

This conclusion becomes less convoluted when we write it in a slightly more symbolic way:

If for some instance of seeming x , it is the case that (for some instance of seeming y , it is the case that P), then for some instance of seeming z , it is the case that P .

I will call this conclusion the *principle of subjective redundancy* (PSR), because it says, in effect, that multiple consecutive occurrences of "it seems that" are redundant.

The PSR has an interesting consequence. It reveals that there is a certain sense in which *you cannot be mistaken about what seems to be the case for you*. This consequence has little to do with older philosophical views about the infallibility of knowledge about one's own mental states (see below for more about these views). The PSR allows that you can make mistakes about what seems to be the case for you; it allows that you might describe appearances wrongly, even to yourself. But it forbids one kind of perceptual error: it says that things cannot seem one way but appear as though they seemed another way. The very idea of such an error is incoherent, as one learns when one fully grasps the notion of subjective fact.

Earlier I pointed out that an immediate past consciousness event can exist for a present consciousness event. When this happens, the present consciousness event and the immediate past one satisfy the PSR, where they play the roles of x and y respectively. Hence one's present perceptions of the

immediate past consciousness event are free of a certain kind of error. But this conclusion holds only for a past consciousness event so recent that it still exists for you, like the consciousness events in examples (1)-(4) above.

The PSR implies that our knowledge of some other consciousness events besides the present one is absolutely trustworthy, though only in a very peculiar and painfully limited way. This absolute trustworthiness can be summarized as follows: If I notice that I just had an experience, then I really had that experience. For all I know, that experience may be deceptive in many respects; there also may be much about it that I did *not* notice. But nevertheless, the experience really happened.

The PSR does not rule out the possibility of mistakes about one's own mental states, or even of mistakes (made slightly after the fact) about the way that things seem to oneself. The PSR appears to be weaker than certain other claims which have been made regarding the certainty of judgments about one's own mental states — for example, the claim of Bertrand Russell,⁵ and possibly that of Peter Carruthers.⁶ The PSR does not address the question of whether it is possible to believe wrongly that one is in a particular mental state; it does not rule out the possibility that I do not feel pain but only believe that I do.⁷ It does rule out the possibility that it seems to me that it seems that I am in pain, whereas actually I do not seem to myself to be in pain. But that is different from being wrong about being in pain. By no means does the PSR imply that *introspection* is

a reliable source of knowledge. It does not even guarantee that introspective knowledge, as ordinarily understood, is possible. Introspection involves sustained observation of, reflection on, and remembering of psychological processes; it does not merely tell us the way that things just seemed for us. The PSR also does not imply that phenomenological descriptions of one's own mental states — descriptions of the sorts used by philosophers of the various schools of phenomenology — always can be trusted.

The peculiar, limited sort of "in corrigibility" which the PSR implies is not threatened by the untrustworthiness of memory. The immediate awareness that something has just happened is different from memory, as memory ordinarily is understood. Instead, it involves something like what Russell called "knowledge by acquaintance."⁸ It happens before memory sets in, so to speak. Psychologists have used the term "sensory memory" to designate the transient sort of memory which persists briefly after a sensation ceases to be actually felt.⁹ Sensory memory may, at least during its earliest phase, involve the existence of an immediate past consciousness event for the present one. However, the existence of an immediate past consciousness event for the present one is not merely a special case of sensory memory, since it can involve an awareness of other mental phenomena besides sensations.

The PSR does allow us to claim qualified infallibility of a very restricted sort for one very special kind of knowledge. This knowledge can best be described as knowledge by acquaintance with what seems to be the case for us. What

we know by acquaintance with our own apparent inner worlds cannot be gainsaid. What seems to be in your inner world, is in fact in your inner world — where "inner world" refers only to the realm of facts that seem to be the case for you. The PSR does not imply even limited infallibility for *all* statements or beliefs about what seems to be the case. (For example, if I try to utter a sentence describing what seems to be the case for me, and this sentence is too long to *seem* to me to have been uttered, then I cannot be absolutely sure what seemed to be the case when I started the sentence.) But if I know, *immediately* after hearing a loud noise, that it seemed that there was a loud noise, then this knowledge is certain. It ceases to be certain very quickly.

The special, short-lived incorrigibility provided by the PSR is intuitively plausible. If I have just felt a great shock, then I cannot reasonably doubt that something has happened, and specifically, that I felt a great shock. I may doubt all sorts of things about the shock: that the shock was objectively real, that there really was a persistent subject who experienced the shock, that the shock really took the amount of time that it seemed to take, and so forth. But the fact that *I felt a shock* is beyond doubt.

(I should mention here that the existence for a consciousness event *x* of another consciousness event *y* does not imply that any particular fact that is the case for *y* also is the case for *x*. The PSR only insures that what seems, during *x*, to be the case for *y*, really is the case for some consciousness event.)

Later I will have more to say about knowledge of

consciousness events. For now the crucial point is the incorrigibility, in a very limited and special sense, of some knowledge of consciousness events other than the present one.

The PSR and Knowledge

The PSR allows one to justify claims about something other than one's present experience, beginning with facts about how things seem now. With its help, one can pass from facts about how things seem to facts about how things seem, or seemed, in another consciousness event. Hence *the PSR enables us to complete one step in the first part of the project of this book*. It does not give us knowledge of reality beyond how things seem, but it does let us draw a conclusion about the reality of the past — or, more precisely, about the status of certain facts which we ordinarily would regard as past. If there is a way things seem, and if it seems that there was an experience which just happened, then we can conclude that there really was a way things *seemed*.

The conclusion that one can know for certain what one has just experienced may seem obvious to everyone except a Humean skeptic.

The PSR probably is the weakest rule that will suffice to get us from facts about how things seem now, to other facts of any sort. Indeed, one can argue that the PSR must hold for knowledge to be possible at all. It is difficult to see how

one can justify any belief grounded in experience unless one grants the PSR, or some principle close to it. If the PSR is denied, something like the following could happen: I feel a great shock, notice that feeling, and yet do not have any grounds for justifying the claim that any experience has occurred. If knowledge is this elusive, then it is difficult to see how anyone can manage to know anything; one could not take note of the immediate past at all, and one could not know what just happened. If the PSR is false, then it is difficult for me even to justify the claim that it *seems to me* as if I just wrote a word. Thus the outright denial of the PSR might well lead to an utterly blank skepticism far deeper than Descartes' proposed doubt.

Fortunately, one never has to live with such skepticism. One can learn that this degree of skepticism is absurd by appealing to the PSR — a rigorized version of the vital logical principle that *to seem to seem, is to seem*.

In the remainder of the book, I will assume that consciousness events other than the present one can be known to exist in the way described in this chapter. Beginning from the experience-derived fact that there are consciousness events which exist for one another, I will attempt to justify other claims about the world.

Chapter 5

Conscious Beings and Their Histories

In Chapter 4 I showed how to take a preliminary step toward the first goal set forth in Chapter 1. To do that, I pointed out a logical fact about consciousness events: that one consciousness event can exist for another. This fact is interesting, not only because of its consequences for knowledge, but because of its bearing on another major philosophical problem: that of *personal identity*. In this chapter I will show how the ideas of subjective fact and of consciousness events can lead us toward a solution to this problem.

Personal Identity: An Introduction

The problem of personal identity¹ is one of the most important philosophical problems from a practical point of view. It amounts to the following question: How do all the different stages and events in a person's life form the life of a single, undivided individual? It is not obvious why these

events and stages don't just exist as separate phases, instead of amounting to the career of one person. If we look at a single snapshot from a person's life — a single moment or brief stage — it may be clear that there is one person there. But if we consider two such stages, perhaps many years apart, what are the grounds for claiming that they really are phases in the career of the same person?

The problem of personal identity becomes acute when we consider that some people change a lot over time, and that all of us change at least a little from moment to moment. The problem asks us to consider what, if anything, remains the same through all these changes.

The philosophical literature contains several different accounts of personal identity. Such accounts examine the conditions under which two given states or stages of personal existence are parts of the career of the same person. I will not attempt here to summarize all of these theories or to criticize them individually. Instead I will refer the reader to the literature on this topic for further information.

Different people have different intuitive views about what must happen if they are to continue existing through time. For example, many people feel that the persistence of memory is necessary for personal survival. On this view, a case of total, irreversible amnesia, followed by relearning of all the facts and skills that one person might know, would lead to the creation of a new person.²

Many philosophers have argued that the continuity of memory, or at least of memory-like mental traces ("quasi-memory"), is necessary for personal identity through time.³

But some people feel that even if they suddenly lost their memories and had to put everything back together from scratch, they still would survive in some form, provided that the "stream of consciousness" (William James' phrase)⁴ containing their experiences is not irreversibly interrupted.⁵ Some philosophers, notably James⁶ and more recently John Foster⁷, have supported views of personal identity in which the continuity of a stream of consciousness plays a central role. Such views differ substantially from those which require continuity of memory. One can think of puzzle cases (usually involving complete forgetting of everything, what Sydney Shoemaker has termed "philosophical amnesia"⁸) in which continuity of consciousness is preserved although continuity of memory is lost. Theories of personal identity also differ from one another in other ways far subtler than the ones I have described here.

Differences among views of personal identity have practical implications, some of them deadly serious. The most dramatic examples of these implications arise in medical ethics. Here I will mention only one such example, based on ones in the literature.⁹ Imagine that a patient has contracted a brain disorder which leads to complete amnesia but not to coma, and which leaves no permanent physiological impairment so that the patient can relearn everything from scratch and thereafter live a nearly normal life. If personal identity depends upon continuity of memory, then the original patient has ceased to exist. Thus, killing the patient immediately after the onset of total

amnesia merely prevents the formation of a new person. Such an act seems at first glance to have roughly the same moral import as contraception; it prevents the creation of an as-yet-nonexistent person. But if personal identity depends upon some version of continuity of consciousness, then the same person likely still exists after amnesia sets in. In that case the killing is a much more serious matter; it is euthanasia at best, murder at worst.

The differences between theories of personal identity sometimes are thought to have important consequences for beliefs about immortality.¹⁰ Suppose that you somehow got the straight information on what will happen to you after your death. Suppose that what you learned was that the perceptual processes now occurring with the help of your brain will either continue somehow in an immaterial soul or be transferred by scientists to the brain of a new body. However, all of your memories (along with "quasi-memories" and the like) of life on Earth will perish with your cortex. Would this form of "immortality" constitute your survival?¹¹ On the continuity-of-consciousness view of personal identity, this scenario may yield real survival — a continuation of your existence, albeit one in which you start all over again as what psychologists call a "blank tablet." On memory-based views of identity, this scenario leaves no hope of survival.

An Agenda

In this chapter I will develop a partial theory of the histories of conscious beings. I will not yet try to pass from facts about how things seem to the conclusion that there are conscious beings which persist through time. (I will address that task in Chapter 10.) But one does not need to assume that there are persisting conscious beings to study those interesting trains of events which we call "histories of conscious beings." For now, one can think of these trains simply as histories of *changing points of view*. Alternatively, one can think of them as *conscious lives* — temporally extended processes involving awareness.

The theory developed here will make use of the apparatus of consciousness events and subjective fact developed in previous chapters. My aim in developing this theory is twofold. First, I want to pave a little more of the road from experience to cosmos by showing that one can infer the existence of a conscious-subject history from facts about how things seem now. Second, I wish to clarify and rigorize some concepts which we often use informally and which will be used more carefully in later chapters of the book. The most important of these concepts is that of *subjective time* — time as experienced by a conscious subject.¹²

Before beginning, I want to examine a more general problem about the notion of personal identity.

The Vagueness of Personal Identity

Philosophers have noticed that the notion of personal identity may be vague in a significant way. Shoemaker has pointed this out explicitly¹³, and also has referred to "a parochial element"¹⁴ present in our usual thinking about that identity. Eli Hirsch has discussed the possibility of alternative notions of personal identity which might appear as normal to some (possible) beings as our notion does to us.¹⁵ The arguments with which these various philosophers support their various conclusions suggest that there is no unique, logically rigorous notion of personal identity, and that our ordinary criteria of personal identity may well contain a conventional (or at least a contingent) element. The differences among different notions of personal identity do make a difference; they can lead to distinct moral and religious conclusions. Hence we must *explicate*, or find a more precise version of, the notion of personal identity before we can hope to compare these alternative conclusions.

My objective here is to define and study one precisification of the notion of personal identity. I will provide a definition of a rigorous notion — that of the identity of a *conscious subject* through time — which corresponds roughly to the notion of the identity of a person. Foster already has proposed an interesting account of the identity of the conscious subject — what he has called "subject identity."¹⁶ My account will be similar to Foster's

in certain respects, though the two accounts differ in important ways.¹⁷ My account of conscious-subject identity is not supposed to capture the entire intuitive notion of personal identity, nor will it agree perfectly with everyone's feelings about personal continuity. (For example, I doubt that every person would feel comforted if it turned out that something identical to him/her in the suggested sense will continue to exist after his/her death — although I think that he/she should feel somewhat relieved.) The notion of the identity of the conscious subject does come close enough to the idea of personal identity to count as one plausible way of making the latter notion precise.

Continuance and Subjective Duration

In what follows I will use the term *conscious subject*, or just *subject*, informally to mean "conscious being." At this stage, I am not yet using the existence of conscious beings as a premise. However, it will be convenient to talk about subjects to motivate certain arguments. Without defining "subject" at this stage, I will take it for granted that a subject is an entity whose history includes consciousness events. This, I believe, would follow if one defined a conscious subject as an entity which is conscious. The most familiar conscious subjects are conscious humans — or, if one prefers, their conscious minds or selves. In Chapter 10 I will take up the topic of conscious subjects again, and will provide a more rigorous characterization of conscious

subjects.

Our immediate aim here is to find out in what the identity through time of a subject's consciousness consists. First we need to find an answer to the following question: Under what conditions do two consciousness events form parts of the same conscious-subject history? This is the analogue, for conscious subjects, of the question of the nature of personal identity.

We can restate the question of conscious subject identity as follows. Consider two consciousness events; call them *x* and *y*. What determines whether *x* and *y* are consciousness events in the same conscious life, or subject history? In other words, how are the consciousness events in the life of a conscious being strung together to form the conscious life of a single being?

In Chapter 4 I discussed the fact that one consciousness event can exist for another. If a consciousness event *y* exists for another consciousness event *x*, then in *x* it seems as if *y* exists. However, in *x*, it may be that *y* does not seem *present*, but seems just past; it may be the case (and normally always is the case) that *y* is not the same consciousness event as *x*. In *x*, it may seem as though *y* just occurred; although *y* seems to be past, some of the subjective content of *y* "carries over" into *x* as part of the realm of subjective fact associated with *x*. From now on I will use the word *continuance* to describe this relationship between two consciousness events. That is, if *x* and *y* are consciousness events and *y* exists for *x*, I will say that *y* undergoes *continuance* in *x*, or simply that *y* is *continued*

during, or in, x.

The next few paragraphs are intended to point out and emphasize some psychological features of continuance. In this paragraph I will speak freely of subjects, experiences, and the like. By doing this, I am not introducing the existence of such entities as a premise. Rather, I am using discourse about such entities to point out certain facts about the way things seem.

Continuance does not occur only during episodes of deliberately focused attention, like those which arise when one works through the examples (1)-(4) in Chapter 4. Continuance occurs all the time during ordinary experience. Normally you do not think about this phenomenon. Yet every moment that you are having experiences, you also experience the fading away of immediately past experiences. For example, continuance occurs when I turn my eyes in the customary way and look at different things. As each new view begins, I "feel," without thinking about it, that what I am looking at has changed. The previous view is no longer seen, but the fact that there was such a view is evident a very brief time after that view ends. A short while later, the previous view fades into memory, or (more often) simply is forgotten.

Immediately after hearing a sudden loud noise, you are aware that something has taken place. The noise still is a matter of "immediate" experience; it has not yet become a mere memory. During the moment immediately after you hear the noise, you are no longer hearing the noise. Nevertheless, you are *immediately, directly* aware that it

happened; the event of its happening still exists for you. At that moment, continuance is occurring. The instance of seeming in which you heard the noise exists for your consciousness, but the noise no longer is heard.

Continuance allows one to be aware that one has just had an experience. Also, it allows one to know this with certainty. These points were made in Chapter 4, where I argued, in effect, that a certain kind of knowledge about consciousness events in continuance is infallible in a limited way. Memory does not share this virtue with continuance. If continuance of a remembered experience is absent, one cannot be absolutely certain, on the grounds of present experience alone, that one has had that remembered experience. There always is the threat of a false memory. But with continuance, such a threat is not an issue. When a consciousness event of yours undergoes continuance, the consciousness event itself exists for you after it ceases to belong to your present experience. The continued experience could not have been pure fantasy, or something implanted in your mind through neurostimulation, as a remembered experience might have been. (If the experience of a continued consciousness event were somehow implanted, then that consciousness event would have to have been implanted also!)

The above remarks reveal a logical connection between continuance and our *awareness of time*. In ordinary human experience, the continuance of a consciousness event makes that event seem to be immediately past, or at least passing. If a consciousness event besides a present one is not being

continued now, then that consciousness event does not appear to be in the immediate past; it may seem to be remembered from the more distant past, or perhaps it does not seem to have happened at all. Hence what is immediately past *for me* — that is, in the time ordering of my experiences as they happen to me — is simply what I am experiencing in continuance.

It is important to recognize that this kind of *psychological* immediate pastness is not the same as immediate pastness in *physical* (clock) time. The difference between these two relations becomes more obvious in cases of anesthesia or very deep sleep. It is my understanding that persons undergoing surgery under general anesthesia sometimes wake up with the feeling that no time has passed since they became unconscious, and that the happenings immediately preceding unconsciousness have "just happened." A similar experience occasionally happens in connection with normal sleep. If an experience of this sort happens, then some final moment of experience, which occurs just before the onset of unconsciousness, must lie in the immediate subjective past of the first consciousness event after awakening. For the subject, nothing has happened in between, although for outside observers time has passed. (Often the subject does not remember the last moments before unconsciousness, but this possibility need not affect the validity of this argument.)

Another example of the difference between subjective and physical pastness comes from certain psychological experiments in which events are perceived to be in the wrong temporal order. Under certain conditions, stimuli

may seem to be in an order different from the order in which the stimuli actually occurred, or it may appear as though later sensations somehow influenced the perception of earlier ones.¹⁸ This suggests that stimuli occurring in a certain order in time may give rise to experiences which occur in the opposite order in the ordering of subjective time provided by continuance. (Of course, there are other possible interpretations of these experiments. Perhaps the experiences occur in the same order as the stimuli, but afterwards *seem* to have occurred in reverse order. This interpretation actually may agree with our first interpretation, especially if Dennett's conception of what happens in these experiments is at least partially correct. On his view, it normally is impossible to say whether the experiences only are recalled as if they occurred initially in the wrong order, or whether they really occurred in that order.¹⁹)

The Stream of Consciousness

A history of a conscious subject can be thought of as the history of a single consciousness as it persists through time. This way of thinking about subject histories is not new; it can be found in Locke's theory of personal identity²⁰ and more recently in Foster's theory.²¹ Using the language of Chapters 2 and 3, we can say that such a history is some sort of series of successive consciousness events, with one event giving way to another. A string of consciousness events of

this sort, with each event giving way to the next, is the only item which can be said to be a process of being conscious — that is, to be the history of an ongoing consciousness.

This view of the history of a conscious subject allows us to form a clearer picture of what holds such histories together. If one consciousness event comes just before another, then the two events form parts of the same subject history. However, it is not important that the second event comes after the first one in "real," physical clock time. It is enough that it seems, during the second event, that the first event just happened. As we have seen, if one consciousness event is continued during a second one, then the first event is in the immediate past, or is entering the immediate past, from the subjective point of view of the second one. This continuance of one consciousness event in the next is what makes one momentary viewpoint "flow into" another to make up the successiveness of our ordinary experience. Hence if one consciousness event is in continuance during another, both events belong to the history of the same conscious subject.

Two consciousness events belong to the same subject history if they are connected by continuance in this way. We can extend this to more than two consciousness events. Suppose that there are three consciousness events a, b, and c, and that a is continued in b and b is continued in c. Since a is continued in b, a and b belong to the same subject history. Similarly, b and c belong to the same subject history. Hence all three consciousness events can be thought of as belonging to the same subject history. In general, two

consciousness events are parts of the same subject history if one can get from one event to the other by tracing a chain of consciousness events, each of which has the previous one in continuance. In such a chain, each consciousness event dies away in subjective time as the next one begins; the new event involves an immediate awareness of the previous event and of some of the content of the previous event. This intimate mingling of consciousness events constitutes the continuity of a *single consciousness* through time. Each event is a moment of experience in the life of that consciousness.

The kind of identity described in the last two paragraphs can be thought of as the identity of a *naked consciousness* through time. (One should remember that it is no more than this. I do not pretend to know whether this kind of identity is the same as personal identity for any sense of "person" richer than "conscious subject" — for example, the moral or legal understandings of a person.)

If one consciousness event "gives way" to another in the manner which I have just described, then the second event can be thought of as a continuation of the same process or "act" of being aware which began with the first event. One can find convenient examples of such continuing "acts" of awareness in one's own life. If you look at something, and then continue to look at the same thing, then the resulting prolonged experience of yours will span many new consciousness events which are connected to the first event in the way I described above. Each consciousness event within this experience (except for the last) is in the

immediate subjective past of another consciousness event within that experience. For all practical purposes, each consciousness event in such a chain embodies the same consciousness as does the previous event. A new consciousness event can comprise a different stage of each of the processes of sensing, thinking, and so forth which began during preceding events.

At each consciousness event in this chain, the relationship between that event and the one before it *seems like a change, or at least like a transition in time*. The following argument explains what I mean by this.

Consider a case in which a consciousness event (call it x) is continued in a second consciousness event, y . There is one point of view, or way things seem, associated with x . There is a different point of view associated with y . Suppose that there is a subject whose history includes x and y . Then both x and y involve pieces of the experience of that subject. However, the subject never experiences both of these instances of seeming as simply being present at once. This is because the subject has no experience of x and y *together*. There is no consciousness event z such that both x and y exist for z . There is no consciousness event z such that all the facts which seem to be the case either at x or at y , seem to be the case at z . Hence a subject cannot experience both x and y as if they were present at once. At any consciousness event, either x seems present, or y seems present, or neither one seems present — but both cannot seem present. Thus, during y , it seems as though the contents of y are there *now*, while the contents of x are not

there now but are close to "now," or are just leaving the "now," or enter somehow into the experience which seems present "now." In other words, from the viewpoint of y, the connection between x and y seems rather like a change.

This argument can be stated less formally as follows. When one sits and stares at a statue, one sees the statue in a continuing way; first one sees it, then one sees it, and sees it, and sees it, and.... Each of these viewpoints involving the statue is a little different from the others; at very least, it involves a sensation or impression or *feel* of having looked a little longer than one had looked during the previous moment. For an experience to persist — to "take up time" or to "last" — is for the experiencer to pass through various slightly different viewpoints in this way. Yet a single viewpoint, by definition, cannot involve passing through various viewpoints in this way. Hence it cannot be felt as something lasting, in the normal sense of "lasting." It does not "go on and on." It must feel as though it were "here and gone" — here during one consciousness event, gone relative to other viewpoints which come after that event in the subject's history.

Thus, when consciousness events are linked by a subject history, their contents must include kinds of experience somewhat like those one normally associates with the passage of time. If one takes "subjective time" to mean the apparent succession of consciousness events along a subject's history, then subjective time *feels like time*. (Of course, many of the features of human time perception — such as long-term memory, expectation of the future, the

sense of time's length, the feeling of inexorability, etc. — may not be common to all possible subject histories.)

The consciousness events in a subject history form what William James called a "stream of consciousness."²² Consider a series of consciousness events connected into a subject history in the way I just described — that is, consciousness events a, b, c, d,... such that a exists for b, b exists for c, c exists for d, and so forth. The event b involves the continuance of a. Thus b is the consciousness event to which a gives way as subjective time passes. The consciousness embodied in b has the event a as part of its subjective realm, so to speak; when b seems present, a seems to die away. A similar continuation of consciousness goes on through c, d,... Each of these events has among its subjective facts the existence of the previous consciousness event. For each consciousness event, the previous "moment" of subjective time is the consciousness event that is just ending. Hence for the consciousness at b, some subjective facts involved in a are in the immediate past. Those subjective facts belong to the fading experiences that happened in the immediate past. We can think of the consciousness in b as a stage in a process of being conscious; the event a which precedes b in the chain also is a stage in this process. It is intuitively plausible to speak this way, because b involves the experiencing as *just past* of some things which for a were *present*. In this way the events a, b, c, d,... make up a single stream of consciousness. Those events are stages in what amounts to an ongoing process of having experiences, embodied at each moment in

some particular consciousness event. Each moment of consciousness in that process "lives on" through continuance into new moments.

The relation of continuance which ties together the stream of experience also provides that stream with an *experienced temporal order*. If a consciousness event *y* exists for a consciousness event *x*, then for *x*, *y* happens "just before" *x*. The event *x* involves continuance, which is a sort of appearance of what has just passed; what has just passed is *y*. Thus we can say that *y* is *immediately subjectively past* for *x* if and only if *y* is continued during *x*. We can define a subjective time order relation in terms of this relation: say that *y* is *subjectively past* for *x* if and only if either *y* is continued during *x* or there is a chain of consciousness events *y, a, b, ..., z, x* such that *y* is continued during *a*, *a* is continued during *b, ..., z* is continued during *x*. (Actually, we only need three consciousness events to construct this chain.) This definition captures what we mean when we say that one experience occurs before another in the stream of consciousness. One cannot plausibly regard a consciousness event of a subject as being past in subjective time unless, in subjective time, it once was *immediately past* — that is, unless one can trace a chain of experience back to the event, by tracing the relation of immediate pastness. Conversely, if an event *x* once was immediately subjectively past (that is, if the event is followed in subjective time by an event, which is followed by an event, ..., which is followed by an event which is immediately past), then it is intuitively correct to say that *x* occurred in the subjective past.

This characterization of subjective pastness does not presuppose the existence of physical time or of physical temporal order. Under ordinary conditions, our experiences unroll as physical time marches forward, but these two time orders are logically distinct. Subjective time order is a felt ordering of experiences; physical time order is established with the aid of clocks or similar physical means. As we have just seen, subjective time order can be defined independently of physical time. Even if it turned out that the physical world were illusory (and I am *not* arguing that it is), there still could be subjective time for conscious beings. The search for a physical explanation for time perception is an important scientific task, but we do not need such an explanation to know that subjective time is real. Whether x is subjectively past for y depends only upon the subjective facts associated with x and with y . (Earlier I mentioned that subjective time may stop while physical time proceeds, if a person becomes unconscious.)

Using this characterization of subjective pastness, we can frame definitions of other subjective temporal notions. For example, by recognizing that a is in the subjective future of b if and only if b is in the subjective past of a , we can obtain a definition of subjective futurity in terms of continuance.

The notion of subjective time discussed above should not be confused with other psychological notions about time. It tells us nothing about phenomena like the awareness of time's apparent length or the understanding of past events.²³ These phenomena are not part of the naked successiveness of experience which I call "subjective temporal order."

An Empty Objection Defeated

One possible objection to the above picture of subject history arises from criticisms of the notion of the stream of consciousness. Dennett, in particular, has questioned this notion. On Dennett's view, the contents of consciousness result from what amounts to the ongoing "editing" of the data of experience, not from one unique, consecutive process.²⁴ But even if Dennett's theory were right, it could not imply that consciousness does not consist of a single stream — provided that we take "consciousness" to mean "the possession of a way things seem" (recall Chapter 2). Even if the stream of consciousness were an illusion of some sort (as Dennett's theory suggests it is), there still would be a way things seem *in the illusion* — that is, there would be subjective facts and consciousness events. Given a particular way things seem, it might sometimes seem that another consciousness event of a particular kind just happened. According to the arguments in Chapter 4, this would mean that there really was such a consciousness event. (This would be the case even if no "conscious" processes had happened in the brain before the later consciousness event. In that case, the "earlier" consciousness event could come into being at the same physical time as the "later" one, yet still be earlier in subjective time.) A chain of consciousness events linked together by this relationship would constitute a subject

history. Hence even if Dennett's "Multiple Drafts model"²⁵ were right, it would not have any bearing on my conclusion that there exist streams of consciousness events, and that the life of a subject consists of a stream of consciousness events.

In my opinion, the stream of consciousness which Dennett's theory criticizes is not the same as the phenomenon which I am calling a "stream of consciousness." The stream of consciousness which Dennett rejects is essentially a series of successive "presentations"; Dennett argues that the presentations which this would require do not really occur.²⁶ The stream of consciousness which I am championing is simply a stream of successive *viewpoints*, whose real nature remains open. The consciousness events in the stream need not be or involve "presentations" of the sort which Dennett rejected. Hence the "stream" presented here is not necessarily the same as the stream which Dennett has criticized. When James investigated the stream of consciousness, I think he had the stream of viewpoints in mind. Note also that the stream of consciousness events need not really be temporally continuous (that is, continuous in physical, clock time); it need only seem continuous. Hence Dennett's objection to the view that consciousness is continuous²⁷ is irrelevant here.

My remarks in Chapter 2 about theories of consciousness are important to remember at this point. No theory of consciousness can force us to believe that there are no consciousness events or that no subjective facts are the case.

At most, such theories can only provide us with views about what those items really are. My account of the stream of consciousness utilizes certain relationships among consciousness events, without regard to what consciousness events really are (material? immaterial? behavioral?). Hence no tenable theory about the real nature of consciousness can contradict my account. Furthermore, my account is not a theory of consciousness and does not imply such a theory. I should mention again that Dennett's theory of consciousness does not attempt to refute subjects' claims about the way things seem.²⁸

Subject Identity During Periods of Unconsciousness

A subject can undergo a temporary lapse of consciousness without starting a new subject history and without any interruption of the flow of subjective time. My earlier remarks on anesthesia should make clear why this is the case. States of total unconsciousness such as deep anesthesia need not interrupt the subjective temporal succession of consciousness events. During ordinary waking consciousness, consciousness events continually transpire as physical time passes. Thus, there is a correspondence between the passage of subjective time and that of physical time. During anesthesia, subjective experience fails to flow during some interval of physical time. But *prima facie*, the stream of consciousness is not

interrupted; instead, the usual relationship between physical and subjective time is modified. Anesthetic states do not really break the stream of consciousness. They merely allow an unusual quantity of physical time to elapse during the transition between one temporal phase of that stream and the next. They also may prevent remembering of subjective facts from consciousness events shortly before the anesthesia.

The above remarks hold for states in which a person becomes totally unconscious — that is, undergoes no consciousness events during an interval of physical time. Most so-called unconscious states are not of this sort. Dreaming sleep is accompanied by some subjective activity and therefore is a segment of the subjective time stream, not a gap in it. Such a condition is not a genuine instance of unconsciousness; it is a condition in which the *content* of consciousness has become markedly altered. The same can be said for any other odd state of awareness in which some subjective life persists. Fugues, near-comas with some residual sensation, periods of what Leibniz called "minute perceptions,"²⁹ and the like do not pose any threat to the identity of the subject. (Whether such states can affect *personal* identity is a separate question.)

Three Technical Notes

In the rest of this chapter I will lay out some technical details of my theory of subject histories. The three technical

notes which follow will be of interest mainly to those with interests in logic or in the philosophy of logic; it is possible to skip these notes without loss of continuity. The first note shows how the ideas of subjective time and of subject history can be made rigorous. It also underscores the point that a conscious subject history is not a logical construction. The second note asks the question "To which ontological category does a subject history belong?" The third note examines some topological properties of subjective time, and some possibilities for unusual topologies of subjective time.

Note 1: How To Formalize Subjective Time

This note indicates how the concept of subject history might be formalized. I will point out one way in which this can be done within a second-order formalized language. (For the required logic and set theory, see texts on those subjects.³⁰)

Let F be a class (or, if one prefers, a property) of consciousness events. Define the *subjective precedence relation* on F as the transitive closure of the continuance relation on F . More precisely, say that a relation R is a subjective precedence relation on F if and only if the following three conditions are met: (1) F is the field of R ; (2) for all x and y in F , if x is continued in y then x bears R to y ; (3) R is transitive on F ; and (4) R has no subrelation besides itself which satisfies (1), (2), and (3). Then define a

subject field as a class F of consciousness events with the following properties: (1) F is *nonbranching* — that is, no x in F is continued by two distinct consciousness events in F or continues two distinct consciousness events in F ; (2) F is *maximal* with respect to continuance — that is, (2a) if x is in F and there is some consciousness event y which continues x , then some such y is in F , and (2b) likewise with "continues" replaced by "is continued by"; (3) F is the field of a subjective precedence relation R on F which is *connected* — that is, for any distinct x and y in F , either x bears R to y or y bears R to x . A subjective precedence relation is what we informally call a relation of subjective pastness or "beforeness." Finally, an object is a *subject history* if and only if it is the *mereological composite*³¹ of all consciousness events in some subject field. In other words, the subject history is the whole of which those consciousness events are parts. The subject history is not the subject field (and hence is *not* merely a logical construct), but is a concrete event or process. It is composed of the consciousness events in the subject field, which can be thought of as its temporal parts in subjective time.

Some readers may be bothered by the idea of a whole whose parts are consciousness events. If consciousness events actually are events, then this whole probably is unproblematical; after all, the consciousness events in a subject history are related to one another in a most intimate way, and usually are spatiotemporally contiguous as well. But in the most general case, consciousness events cannot be supposed to be events; all we know for sure is that they are

instances of seeming. It is difficult to imagine how instances of seeming which are not events could be the parts of a whole! I will address these difficulties in the next technical note.

The definition of subject history, whether in the rigorous form above or in the informal version given earlier, reveals the following important properties of subject histories. A subject history is a single stream of consciousness; it cannot be, for example, two parallel streams of consciousness, or a swarm of disconnected consciousness events. The connectedness condition on the subjective precedence relation R insures this uniqueness of the stream. The stipulation that the subject field is nonbranching implies that for each consciousness event x in the subject history, there is a unique, linearly ordered series of consciousness events in the history which lie in the *near* subjective past and future of x (provided that x has a subjective past and future). That is, some segment of subjective time around x has a linear topology. The maximality condition on a subject field insures two things: (a) if a consciousness event x in the history of a subject gives way to some consciousness event y (that is, if x is continued during some y), then some such y also is a part of the history of that subject; (b) if a consciousness event y in the history of a subject has some consciousness event x in continuance, then some such x is part of the history of the same subject as y . In other words, the subject history does not begin later than, or end earlier than, the stream of consciousness. Hence any consciousness event which is part of the same nonbranching "stream of

consciousness" as an event x will belong to the same subject history (or histories) to which x belong(s).

This definition of subject history captures the informal notion of subject identity which I explained informally above. Intuitively, two consciousness events are events in the career of the same subject if and only if they belong to the same subject history.

Note 2: The Ontology of Subject Histories

Intuitively, one may think of a subject history as an event — specifically, as a temporally extended event which has consciousness events as parts. If a consciousness event is indeed an event, then my definition of a subject history agrees with this intuition. However, there is no *a priori* guarantee that all consciousness events really are *events* in the usual sense, or are items that happen in physical time. Thus, we cannot rule out subject histories which are not events or which do not occur in physical time. However, we are safe in regarding a subject history as a certain kind of whole having consciousness events as parts. If the consciousness events really are events, then the history is an event.

A further question arises when we consider the whole which the consciousness events are supposed to form. If consciousness events really are events, then it is possible to assume that these events form a whole, especially in view of the intimate way in which the events are interconnected. This plausibility increases if the events are, for the most

part, contiguous in time and space — as neural events in a single brain might be. It is likely that all consciousness events are events, so a whole composed of consciousness events probably is no more problematical than any other events composed of multiple temporal parts. However, we have not assumed that consciousness events are events. Would instances of seeming which are *not* events form wholes in the required way?

The answer to this question is implicit in the definition of consciousness events as instances of seeming. In Chapter 2, I pointed out that the existence of an instance of seeming or consciousness event does not involve anything over and above facts about how things seem. There is nothing more to the existence of a consciousness event than the obtaining of certain subjective facts. A similar statement can be made about wholes composed of consciousness events. The claim that there is a subject history says nothing more about the world than does the claim that consciousness events of certain sorts exist. (The required sorts of consciousness events include consciousness events for which other consciousness events exist, and which are ordered by this interrelationship in a certain specific way.) This last claim, in turn, says nothing more about the world than does the claim that things seem certain ways in certain instances. Thus, the claim that there are subject histories is as secure as the claim that things seem certain ways. The ways things must seem to make a subject history exist are rather specific; certain instances of seeming must seem in certain other instances to exist, as detailed in the definition of a subject

history.

Note that we may interpret quantifiers over subject histories substitutionally, as we did for consciousness events (and for the same reason).

Those who truly detest the idea that instances of seeming form wholes are free to adopt some other view of what a subject history really is. For example, one might think of a subject history as a *property* of consciousness events. All the consciousness events in John's subject history could be regarded as possessing a common property — say, that of being a "John-consciousness event." One could just as well regard John's subject history as a class of consciousness events (that is, identify the history with its subject field). One might even think of a subject history as a state of affairs involving consciousness events. For example, one can take the real content of "there is a John-history" to be the fact that there are John-consciousness events and non-John-consciousness events in the world. All of these alternatives, particularly the one involving classes, amount to the use of logical constructions as subject histories. As I said earlier, my aim in this book is *not* to find logical constructions which will substitute for objects, but to learn something about the objects themselves.³² I mention these three alternatives, not because I advocate them, but because they allow those who reject my characterization of subject histories to continue reading the book. One can accept much of what comes later in this book without believing that subject histories are wholes made of instances of seeming.

One might wonder whether subject histories even need to

fit into any of the standard ontological categories. Entities as special as *changing viewpoints* or *streams of seeming* might not exactly fit under any other heading. Perhaps subject histories are just — subject histories!

Note 3: The Topology of Subjective Time

The relations of continuance and of subjective pastness have certain formal properties which possess clear psychological meanings. Some of these properties follow from the definitions of continuance and of subjective pastness; others cannot be obtained deductively, but are suggested by ordinary experience. Here I will review some of these properties very briefly.³³ This note presupposes a knowledge of the elementary theory of order, such as is discussed in texts on set theory.

Reflexivity. In ordinary experience, the relation of continuance is irreflexive; a consciousness event does not "contain" itself in the way in which a consciousness event "contains" another consciousness event in continuance. However, the definition of continuance offers no prima facie guarantee of this. Similarly, in ordinary experience subjective pastness is irreflexive; a consciousness event is not experienced later than itself. But the definition of subjective pastness does not guarantee this. Also, the irreflexivity of continuance does not imply the irreflexivity of subjective pastness.

Symmetry. Reflection on everyday experience suggests that continuance is antisymmetric. A human being normally does not have a consciousness event a, then have another consciousness event b in which a is continued, and then immediately have a again. However, the definition of continuance does not, *prima facie*, exclude this possibility. Also, we have no *a priori* guarantees that subjective pastness is antisymmetric. The antisymmetry of continuance does not imply the antisymmetry of subjective pastness. If subjective pastness failed to be antisymmetric, then there could be consciousness events x and y such that x is both before and after y in subjective time. This would happen if the topology of the subjective time of a subject were closed.³⁴ The irreflexivity of subjective pastness also could fail under these conditions. Such things might happen to a physical observer in a universe which has closed time or permits time travel. I do not know of a way to rule out this possibility.

Transitivity. In our ordinary experience, continuance is not transitive. If it were, then a subject could, at any moment in his/her history, notice all of his/her past experiences in continuance. Such a subject would experience his/her entire past as immediately past; that entire past would seem that it had "just happened." If a subject history (as I have defined it) were like this and also contained more than two consciousness events, then there would be branches in the subject history (a distinct

consciousness event would have to continue more than one consciousness event). The definition of subject history rules this out. Hence there cannot be a subject history (as defined above) in which continuance is transitive, unless there is a subject whose history contains only two consciousness events. For such a short-lived subject, continuance would be vacuously transitive.

Trichotomy and nonbranching. In ordinary experience, continuance does not obey the trichotomy law on consciousness events in a subject history. If x and y are consciousness events in the same subject history and x is in the distant subjective past of y , then x is not continued in y , y is not continued in x , and y is not identical to x . Subjective pastness, restricted to a single subject history, obeys the trichotomy law.

The trichotomy of subjective pastness is an important feature of subjective time. Subjective pastness in a subject history obeys a trichotomy law: for consciousness events x and y , either x subjectively precedes y , or y subjectively precedes x , or x is y . (Since we have not ruled out universes with closed time, we cannot generally regard these "or's" as exclusive.) This trichotomy law excludes cases in which two or more streams of consciousness are parts of the history of the same subject. For example, if a subject splits to give two subjects, the resulting pair of streams of consciousness do not make up the history of a single subject. (I will discuss puzzles about splitting and merging subject histories in Chapter 12.)

For partial orders, trichotomy is known to imply the absence of branches in the order. This implication does not hold in general for the subjective pastness relation; since we cannot rule out the possibility that this relation is topologically closed, we cannot be sure that it is a partial order. Hence a separate nonbranching condition is needed in the formal definition of subject history (recall Appendix A).

Local properties. Subjective pastness also has a significant local topological property: for a subject history with more than two consciousness events (or for any subject history which is not closed), the subjective pastness relation is a linear order when restricted to a sufficiently short segment of the subject history. This is a direct consequence of the definition of subject history.

Summing up: By using the definitions presented in this chapter, we can show that the subjective pastness relation is transitive, and is trichotomous (in a nonexclusive way) if restricted to a single subject history. Ordinary human experience suggests that for human consciousness under ordinary conditions, continuance is antisymmetric and does not obey trichotomy, and that subjective pastness is irreflexive and antisymmetric. For subjects having three or more consciousness events, continuance is not transitive. Subjective pastness behaves like a linear ordering over sufficiently short stretches of an ordinary (that is, non-closed) subject history.

From Brain to Cosmos

Chapter 6

Knowledge of Other Minds

It is a fact of everyday life that one person cannot directly witness another person's experiences. Philosophers who think about this fact have encapsulated it in the commonly made claim that experience is *private*. The inability of persons to witness the experience of other persons is, at very least, an important part of what philosophers have meant by the privacy of experience.

Philosophers have responded to the apparent privacy of psychological life in a variety of ways. Some thinkers, notably dualists, have given this privacy great significance.¹ Others, most notably the behaviorists, have tried to deny the existence of private mental processes.² The most important and best known problem raised by the apparent privacy of mental life is the *problem of other minds*.³

The problem of other minds may be stated as follows. Granted that you cannot directly witness other people's mental processes, how can you know that other people have any mental processes at all? Even if you are sure that you do

know this, there is a puzzle about *how* you know it. Imagine that you lived in a world in which you were the only conscious self, and other so-called persons acted exactly like real persons but lacked consciousness. In such a world, your experiences of other persons would be exactly the same as they are now. Since all of your knowledge about other people's thoughts, feelings, and the like is based on your experiences of other people's bodies and behavior, how do you manage to know that other people really have minds? How do you know that they don't just *act* like they have minds?

We also can state the problem in a less dramatic (and more general) way as follows: How can I know what is going on in another person's mind? What are the criteria for inferring that someone else is undergoing a conscious experience?

In this chapter I will investigate the problem of other minds with the help of the ideas developed in the preceding chapters. I will argue that under certain conditions, one can infer that there is a subject other than oneself from facts about how things seem to oneself.

The arguments in this chapter make use of the conclusions drawn in earlier chapters. However, these arguments are, for the most part, non-rigorous. Some of them are meant to motivate or illustrate certain concepts rather than to establish conclusions. At one point I will use an argument based on facts from psychology and biology to support one of my conclusions. The premises used in this argument go far beyond facts about how things seem. I have

included this argument to illustrate and lend plausibility to my thesis, and to show that my thesis does not conflict with certain widely held views derived from science. (Because of its non-rigorous character, this argument, and much of the rest of the chapter, belongs to the second part of the project of the book, as described in Chapter 1.)

The Privacy of the Psychological

In Chapter 5 I proposed a criterion for conscious subject identity. There I said that two consciousness events are part of the same subject history if they are connected by a chain of consciousness events, each member of which involves continuance of the previous member of the chain. This stipulation has an interesting consequence: a subject ordinarily cannot experience in continuance a consciousness event which belongs to another subject's history. The following argument shows this. If subject John experiences the continuance of a consciousness event which is in subject Jack's history, then there has to be a consciousness event y in John's history during which this experiencing occurs. The consciousness event in Jack's history which John experiences in continuance may be called x . Since x exists for y , x is in the immediate subjective past of y . It follows that every consciousness event in the subjective past of x also is in the subjective past of y . Therefore, up to the subjective moment x , John and Jack share exactly the same past! This cannot happen if John and Jack are two different

subjects. The only exception is if John and Jack were created by the splitting of a single conscious subject into two subjects. But ordinary subjects do not originate in this way and therefore cannot hold each other's consciousness events in continuance. (I will have more to say about splitting subjects in Chapter 12; for now I will only mention that there is a philosophical literature on splitting subjects.)

The preceding argument shows that a certain sort of privacy for mental processes follows from the structure of subject histories. A subject's inner contents are not *directly* accessible to another subject's awareness, except perhaps in odd cases in which subjects split. The privacy of the psychological, in this restricted sense, does not involve any mystery. Even if a subject's experiences could somehow be observed by others (for example, if behaviorism were true), those experiences still would not be *undergone* by others. The experiences would not be lived through by external observers in the way that they are lived through by the subject.

Perception of Other Subjects' Subjective Content

We can reformulate the problem of other minds in terms of consciousness events as follows: How can a subject know what is the case for a consciousness event in a subject history not his/her/its own? This question does not capture the entire content of the problem of other minds, but it

captures the most challenging part. I will not try to analyze this entire question in terms of subjective fact, since I have not framed a definition of knowledge in terms of this concept and do not intend to do so. Instead, I will sidestep all questions about the nature of knowledge and will simply try to find a way to determine what is the case for a subject who is not me.

In Chapter 4 I argued that we sometimes can know for certain whether a fact is true for a consciousness event in the immediate subjective past. Clearly this account cannot be extended to any arbitrary consciousness event. Consider a consciousness event in your own distant past. You cannot be certain what facts were the case for this event; your knowledge of those facts will rest on memory and perhaps on non-deductive inference from present facts, and both of these sources of knowledge are fallible. This fallibility becomes particularly serious for consciousness events which are not part of your history at all. What if the consciousness event belongs to another person, and you wish to know about some secret thought which that person harbors? In this case you cannot know immediately what is the case for that consciousness event, and memory is no help either. What other sources of knowledge could be of use?

In real life we garner information about other persons from our observations of those persons' bodies. This information comes by way of our own sense experiences. We know something about other persons because certain bodily facts about those persons are the case for us. In particular, we learn something about persons' mental

contents from those persons' behavior. But philosophers have noticed⁴ that this method of knowing cannot be reliable unless there is a *dependable correlation* between mental contents and observable behaviors. In my terminology, this would amount to a dependable correlation between the contents of consciousness events in one person (the observed) and the contents of consciousness events in another person (the observer). An observer may have the sorts of experiences which we call experiences of another person's behaviors. If the required dependable correlation exists, then the observer can get information about the mental states of the observed person.

Consciousness events which are not in one's own history do not exist for one's own consciousness events. But this does not rule out the possibility that *general facts* about external consciousness events may be the case for one's own consciousness events. Perhaps you could become aware that some fact is true of consciousness events in another subject, even though the consciousness events themselves do not exist for you.

Reflection on ordinary experience discloses certain happenings which appear to involve knowledge of this sort. I am referring to one's everyday "instinctive" or "gut" perceptions about the psychological states of other persons. Some psychologists and philosophers have recognized that perceptions of this sort occur and can convey information. Such perceptions are quite normal; they do not involve anything like mind-reading. Consider the fact that one often can notice when another person is afraid. I am referring to

the "gut reaction" one has to the presence of fear in another person, not to any inference based upon descriptions of that person's behavior. This reaction often occurs before one has time to think about the observed person or the objective situation. It does not require any conscious logical thought on one's own part.

Biological evidence suggests that the direct communication of information through emotion is a phenomenon common in mammals, both human and nonhuman.⁵ Scientists have proposed that one biological function of emotional expression is communication.⁶ The reception of such communications would form an example of "gut" perception of the sort I have described.

Philosophers have uncovered other possibilities for access to others' mental states. On P.F. Strawson's view of persons, the attribution of mental states to other beings is a prerequisite for the use of certain kinds of mental language.⁷ I take this to imply that our knowledge of others' mental states is not a matter of mere inferences from facts about behavior. Frank A. Tillman has studied (from a phenomenological standpoint) the idea that one may notice conscious states in other persons.⁸ I will have more to say about Tillman's ideas later.

Emotion supplies the clearest examples of direct perception of others' psychological states, but the same sort of perception evidently occurs with other psychological phenomena as well. For example, one sometimes feels that one can "just see" that another person is thinking or

concentrating very hard.

It is important to remember that perceptions like these require only the normal processing of sensory information. This point bears repeating because some people may find it implausible that the "direct," perceptual knowledge of other minds could occur through normal sensory perception. Yet such knowledge can be understood neurophysiologically and does not require anything like mind-reading. The fact that the acquisition of such knowledge is not simply a matter of sensation does not make this knowledge any less "direct" or immediate. Other forms of perception also have neural mechanisms, but that does not make them any less direct. Of course, we cannot yet use ideas about neurophysiological mechanisms in the deductive argument from facts about how things seem, since facts about neurophysiology go beyond such facts. The last few paragraphs belong to the second part of this book's project (as described in Chapter 1). They are intended to point out that perceptual knowledge of the mental states of others need not involve anything paranormal — or even anything unusual.

If one notices that another person is afraid, overjoyed, or deep in thought, one is noticing a psychological fact about that person. Then it is the case for one's consciousness events that fear, joy, or some other psychological phenomenon is occurring. But this does not imply that one actually notices the consciousness events occurring in the other person. For despite our ability to be aware that another person is afraid, we still cannot notice that person's consciousness events, even though those are associated with

subjective facts involved in the person's feeling of fear. Another, rather loose, way to put this is as follows: you can notice that someone is afraid, but you cannot notice that person's *experience of fear*.

The preceding distinction is rather subtle, but it makes a big difference. To be aware that Jack is afraid is to be aware that a certain fact about Jack, or about Jack's subjective realm, is the case. If John notices that Jack is afraid, then it is the case for at least one of John's consciousness events that Jack is afraid. But to be aware of someone's experience of fear is to notice instances of seeming; if John actually notices Jack's experience of fear, then certain consciousness events of Jack's also must exist for John's consciousness events. A subject who does not share Jack's past cannot experience Jack's consciousness events in continuance. Hence that subject cannot literally be conscious of Jack's experience of fear. Nevertheless, such a subject might be conscious of the fact that Jack is afraid. Thus the privacy of consciousness events, in the limited sense described above, does not logically exclude the possibility of perception of the psychological states of others.

The Logic of Noticing Other Subjects' Mental States

The claim that one can notice that a person is afraid without noticing that person's *experience of fear* may seem strange. Actually, there is nothing strange or obscure about

this distinction. People often notice that something is the case without noticing the events whose occurrence makes it the case. For example, you can notice that a wall is brightly lit without noticing events of reflection of photons by the wall. Similarly, it is possible to notice that someone is experiencing fear without noticing the consciousness events which play parts in that experience of fear. To notice that Jack is afraid is to notice that a certain fact is the case. To notice Jack's fear (or Jack's experience of fear) is to notice events of a certain sort — either Jack's experiences, or his consciousness events (which either are, or are much like, real events).

The possibility of noticing that someone fears without noticing the event of their fearing is an example of the logical incompleteness of consciousness events, which I mentioned in Chapter 3. This incompleteness implies that even if a fact P implies a fact Q, and P is the case for a consciousness event x, it does not have to follow that Q is the case for x. The fact (P) that Jack is experiencing fear implies the fact (Q) that Jack has consciousness events which involve fear. If John notices that Jack is afraid, then it is the case, for some consciousness event of John's, that (P) Jack is experiencing fear. However, this does not imply that for some consciousness event of John's, (Q) there are consciousness events of Jack's which involve fear. John need not experience Jack's consciousness events.

Similar failures of logical completeness may occur during experiences of illness. Sick people typically become aware that they do not feel well without being able to see or

otherwise directly experience the cellular causes of their illnesses. Suppose that you contract a cold and begin to feel ill. You are aware that you feel a certain way; the fact that you are having sensations of a particular sort might, for all we know, even imply that you are suffering from a cold. Yet you are not immediately aware of the cold viruses themselves; without expensive instrumentation, you cannot even see these viruses. In this example, you notice that you feel a certain way, but you do not notice the virus, even if the fact that you feel that way implies that the virus is present. It is the case for my consciousness that I feel a certain way, and perhaps it even is objectively true that if I feel that way then viruses exist. But it is not the case for my consciousness that viruses exist.

The Fallibility of Perception of Other Minds

The perception of mental states in other beings is quite fallible. Philosophers have framed arguments in which an actor puts on a very good imitation of pain, which cannot be distinguished from the behavioral correlates of real pain.⁹ It is conceivable that I might see and hear such an actor and have the same "gut reaction," or subjective emotion-laden perception, that I would have in the presence of real pain. To what extent do examples like this cast doubt on the reliability of our perceptions of facts about other minds? How can I know that an experience of mine which seems to be an experience of a mental state external to myself is not

simply an experience of a well-done piece of fakery?

The following partial answer to this question is compatible with the account of conscious subjects outlined in Chapter 5. This answer can be thought of as a variant of Frank Tillman's account of the perception of other selves. Tillman suggested that by "reducing the ambiguity" in a certain manner in our experiencing of certain perceived behaviors, we come to know, non-inductively, of others' mental states.¹⁰ Here I will defend a version of this answer, with my own changes of detail.

Suppose that an actor appears on the street before me and suddenly feigns terror. I may have the same "gut reaction" that I would have to real fear. Of course, I am not actually noticing that someone is afraid; I am only seeing a form of deliberate physical activity (or excitement) which the actor performs in order to create an impression of fear in others. My gut reactions cannot distinguish immediately between this imitation and real fear. But after concentrating carefully on the actor's movements and expressions, I might finally notice something out of place. I might notice that the actor seems to be *concentrating on his appearance* while acting — that he seems to be "keeping an eye on himself" — or that he is making an effort to act well. These other perceptions (or perhaps I should call them *suspicious* of mine?) would make me feel unsure about what sort of state I really am seeing.

In this example, my initial perception of emotion was based upon incomplete information. My mind had processed some sensory information and perceived "fear."

But after seeing more, I began to undergo new perceptions which altered that perception of fear. The key fact here is that my initial perception of "fear" was based on insufficient information. I misperceived the actor's emotional state because I had not had the opportunity to notice all the relevant sensory cues — much as I might misperceive a textbook optical illusion before I examine the diagram very closely.

It is worth noting that in the above scenario I am not wrong about the actor's being *in some mental state or other*. I was only wrong about the *content* of that state — about *what kind of* mental state the actor was in. I noticed that the state was one of arousal or excitement, but I did not correctly perceive the subjective content of the state of arousal into which the actor had passed; it was a state of *intense mental concentration* (of the sort required for convincing acting), not one of *fear*. Upon getting a less one-sided impression of the actor's bodily state, I noticed more accurately some characteristics of his mental state. To do this I had to make a mental effort — my own effort of concentration.

This example illustrates why your initial impression of someone's mental state can be misleading. It can be deceptive because of *what you fail to notice*. In the above example, after collecting enough impressions of the actor's actions, I finally notice the actor's mental state: one of intense, deliberate concentration.

It should be clear from the above example of the actor that the direct perception of psychological states is fallible

but correctable, and is reliable only to the extent that it is *unambiguous*. Thus we arrive at an idea close to Tillman's idea which I mentioned earlier.¹¹ The main difference between my proposal and Tillman's is that on Tillman's account, a certain lessening of ambiguity makes the perception of another's mental states inevitable,¹² while on my account, the lessening of ambiguity is of a slightly different kind and makes the *content* of a mental state more clearly discernible.

Trustworthy perception of another subjects' mental state is possible. To be completely trustworthy, such a perception would have to be founded on subjective facts which leave one no choice as to what the state is. All examples of misperception of mental states must violate this requirement to some extent; that is, they must contain some perceptual ambiguity. Such misperceptions involve perception of a mental condition based on subjective facts which together do not contain enough information to specify just what the state is. A particular perceived mental state could be fear or feigned fear — or perhaps even joy in a person who, because of neurological abnormalities, reacts in an unusual way when possessed by joy. But if one does not perceive with the help of enough subjective facts to disclose the nature of the mental state, one might notice only that there is a mental state — if one even notices that much.

Despite its limitations, perceptual access to mental situations external to oneself provides a knowledge of other minds far more secure than anything that can be obtained through rational inference. This is true even if such

perceptual knowledge is quite unreliable. If a link between mental states and bodily states of other beings is only something that I infer, then this inference is quite uncertain. Unless I base this inference on a particular philosophical view of the nature of consciousness, the inference must be inductive rather than deductive.¹³ Such an inference is well known to require generalization from facts about my own consciousness to conclusions about consciousness in general.¹⁴ This generalization from the consciousness of *one subject* to all consciousness everywhere is known to be a rather weak inference — how do I know that it isn't just my own mind that works that way?¹⁵

Summing Up

The above remarks on perception of other minds lead up to my proposed partial solution to the problem of other minds. I summarize this solution (and its limitations) as follows.

A subject cannot experience in continuance the consciousness events in the history of another subject. In this sense, a subject's mental states are private. When we say that people's thoughts and feelings are private or personal, this is the kind of privacy we primarily have in mind — no one else actually can share our experiences. However, there is another way to perceive mental facts about other beings: some such facts can be true for the observer's own consciousness events. Ordinarily, we do not

distinguish between these two kinds of perception of others' mental states, but the difference is important. Perception of the second sort is grounded in sensations of others' bodily states. For it, mental facts about others are not "private." This capacity to perceive the mental condition of others is one example of a deep property of consciousness events: their logical incompleteness.

The ability to perceive others' subjective states is quite fallible; it does not provide a way to find out, once and for all, whether one's perceptions of the mental states of another being are right. However, the knowledge obtained from such perceptions is *self-correcting*. Errors in perception of others' mental states result from perceptual ambiguity; one cannot be sure one has perceived accurately unless one's experience of the physical state of the other subject determines the other subject's mental state uniquely. If one's experience of another subject is rich enough to meet this condition, then the associated perception of mental states is reliable. If one's experience does not contain enough information to specify fully the character of the mental state, then the impression of the content of the mental state can be a red herring.

Chapter 7

The Flow of Time

In this chapter I will examine one of the most important and puzzling features of the world which we experience. This feature is the flow or movement of time. The phenomenon of time's flow, or *temporal flux* as it often is called, is the subject of some philosophical problems. The greatest riddle about the flow of time is the question "Is it real or merely apparent?" Near the end of this chapter, I will propose an answer to this question. I will conclude that the flow of time is indeed real, although it is not what many philosophers of time suppose it to be.

The notion of temporal flux occupies an interesting place in the history of philosophy. Some philosophers have held that the flow of time is an objective feature of the world. According to this view, the present moment really does move, in some sense, from the past toward the future, and the existence of past and future things (if such things exist at all) may differ somehow from that of present things. Process philosophers even have held that temporal flux, or

something closely akin to it, is the very foundation of reality.¹ On the other hand, some philosophers have denied the reality of the flow of time. These thinkers have argued that the apparent movement of time is (in one way or another) an artifact of our perceptions; often they have tried to reduce this movement to relationships among tenselessly existing events or to some other nonmoving substratum.²

Throughout this chapter I will make free use of examples about physical objects to support my conclusions about time. One might think that the facts used in these examples go beyond facts about how things seem. However, they usually do not. Most of the descriptions of physical objects in this chapter actually are descriptions of how physical objects appear to observers; hence these examples actually are examples of how things seem. Physical object examples which cannot be read in this way belong to the second part of the project of this book; those which can be read as arguments about how things seem may belong to the first part. In a similar spirit, I will make use of some facts from physics to buttress certain steps in the argument.

Temporal Flux and Bergsonian Duration

To kick off this discussion of temporal flux I will examine the concept of *duration* as set forth by the French philosopher Henri Bergson. My aim in doing this is not to champion Bergson's philosophy as a whole, but to better understand the notion of temporal flux. Bergson's view of

time, to which this chapter owes much, is the prime example of a philosophical view on which temporal flux is real and irreducible. According to Bergson's view, temporal flux is something distinct from, and not reducible to, the temporal ordering of events.

In ordinary usage, the word "duration" refers to length of time, and secondarily to the persistence of objects through time. According to the latter usage, an object endures if it exists at each time during some interval of time. Some philosophers use "duration" in this way. But Bergson used "duration" to refer to a special aspect of time which cannot be reduced to the relations of order and temporal distance which hold among instants and events.³ This special notion of duration is the one I will examine here.

Bergson understood duration as a sort of pressing forward of the present toward the future.⁴ The Bergsonian conception of duration cannot be explained in a few sentences; the reader is referred to Bergson's works for a full account. The example of the next paragraph, loosely based on one of Bergson's,⁵ illustrates one particular aspect of this conception.

Consider what happens to a coffee cup between 12 noon and 12:02 pm on a given day. Suppose that the properties of the cup do not change during the interval from 12:00 to 12:02 — that is, nothing happens to the cup during that interval except that the cup continues to exist. Consider the part of the career or history of the cup which extends from 12:00 to 12:01. (This part is what sometimes is called a

"thing-stage."⁶) For convenience, let us call this part A. Now compare A to the longer part of the cup's career lasting all the way from 12:00 to 12:02. Call this longer part B; note that A is a temporal part of B. During B, as during A, the properties of the cup do not change; seemingly, nothing happens to the cup during either A or B. But further reflection reveals that something does happen to the cup during B that does not happen during A: after A ends, the cup *exists some more*. As the cup persists through time, it exists, and then exists some more, and then exists yet more. This is what is "happening to" the cup during the entire length of B. This "existing and existing" is what we may call the *enduring* of the cup. It is something that goes on even in the absence of changes in the cup.⁷

This example illustrates a fact about duration which is implicit in Bergson's thought but which a naive understanding of duration may overlook. This is the fact that if an object is persisting through time, *something is happening*. To see this "something" happen, just pay attention to any physical object. As one watches the object, the object persists; as one continues to watch the object, it persists, persists, and persists some more, even if one does not witness any change in the object, and even if the object in fact does not change. This fact of things' "persisting and persisting" is one aspect of what Bergson called "duration." The kind of duration which I have just described is a feature of a thing which is different from the mere existence of that thing through an interval of time. It is something that goes

on at *each* instant of time; at each moment in its history, an object is busy enduring. Duration of this sort is a simple flowing-onward of things. Like Bergsonian duration, this duration is distinct from any time-ordering of the stages of a thing's existence (it is not a mere relation among those stages), and from any metrical, or distance, properties of time (it is not a mere measure of time interval).⁸

The above informal remarks are not intended as a precise definition of duration or as an argument for the reality of duration of the Bergsonian sort. I must stress that they do not do full justice to Bergson's rich idea of duration, although they do capture one side or facet of that idea. They are meant only to paint a portrait of duration as it appears in everyday experience. Now I will attempt to make this informal notion rigorous.

Think of the coffee cup again. Consider a brief slice of the history of the cup, centered at 12 noon on a particular day. Take the slice to be instantaneous or of very short length. In philosophical terminology, this slice is called an *object-stage*.⁹ Call this object-stage S.

According to what I said before, duration is something by virtue of which an object, as it is now, plunges forward toward the future.¹⁰ Hence duration must be something present at each stage in the career of an object. One might say that duration is a feature which belongs to the state of an object, and which ensures that the object will go into other object states at later times. (This description will be useful later, when I will look at the concept of duration in a new

way — one which Bergson probably did not foresee.)

What does the cup's possession of duration imply? At very least, we can say this: Because the cup has the property of duration at 1 pm, the cup at 1 pm *already is in transition* from S to subsequent object stages. At the time when S becomes the cup's present stage, S already is giving way to a future object-stage of the cup. An attempt to examine the cup while it is in a stage S does not reveal a static moment in the history of a cup. Instead, it reveals a view of the cup already passing from S into subsequent object stages. Metaphorically speaking, we can say that the cup *refuses to sit still* at the stage S.

This fact that the cup "refuses to sit still" in time is at least part of what we mean when we say that the cup "endures" — at least if we understand duration in roughly the same way Bergson did. While the cup is at stage S, the cup is enduring; it is enduring because the fact that it is at S implies that it is going to be at other stages at later times. The fact that S is going to lead to other stages in this way is a feature of S itself. When the cup is at the stage S, something is happening to the cup that will take the cup out of S and into other stages. Speaking loosely once more, we can say that the cup endures because each object-stage in its history has the property of giving way to other object-stages. That is, an object endures if its object-stages are transitory *by their very nature* — not merely "transitory" in the conventional sense of occupying a short stretch of time.

Thus, the transitoriness of an object-stage of S consists in a certain kind of implication of other object-stages. The fact

that the object-stage S exists implies that future object-stages exist; hence the object to which the stage S belongs is, in a sense, pressing forward into the future. Of course, the object may cease to exist, in which case some final object-stage does not meet this criterion; but that final object-stage still is transitory, because its existence follows from that of some other object-stage and is followed by nothing (at least by no object-stage of the same object).

Our experience of this transitory character of momentary states lies at the core of the intuitive feeling that time flows. The fact that the present situation gives way to future situations is what leads us to believe that time really "moves." This transitory character of the moment is not the only feature of the world which might be called the "flow of time." In particular, the fact that events seem to "move" toward the past from the future, becoming present for an instant in between, makes time resemble a flow.¹¹ But one does not need to perceive this last kind of "movement" to have an impression of the ephemerality of the present or of the "flowing" character of time. One has only to look at the present moment in the right way. The fact that the present moment is yielding to another moment is what makes time "fly."

The apparent movement of events differs in significant ways from the kind of duration described above. In particular, a flow of events toward the past would depend upon changes in the locations in time of those events, but "duration" as I have described it is a feature of an object at a *single* time. I should add also that D.H. Mellor has provided

an analysis of the apparent movement of events within a tenseless theory of time.¹² If one believes this analysis (and I think it is at least plausible that some analysis of this general sort can work), then the apparent movement of events can occur even in a world lacking what I call temporal flux. (Note that I am not endorsing all of Mellor's views on tense here, only his conclusion that time can seem to flow in a tenseless world. Elsewhere I will endorse another of Mellor's views about time.)

The following quasi-formal definition of temporal flux is motivated by the above remarks.

Let X be an object and t a time. X is *in temporal flux* at t if and only if either:

(1) at t , X is in a state S for which the following condition holds: that X is now in S necessarily implies that at some time t' later than t , X will exist;

or

(2) there is a time t'' earlier than t such that at t'' , X is in a state S such that the fact that X is in S at t'' necessarily implies that X exists at t , and t is the last time in the history of X .

Note that I used "necessarily implies" in this definition, instead of just "implies." The appropriate notion of

necessity here is what philosophers call physical necessity. That is, it must be physically impossible that the antecedent holds but the consequent fails.

The times referred to in this definition may be taken to be the physical times measured by some standard clock, or they may be taken to be moments of subjective time — that is, consciousness events in some subject history. Hence we can speak of temporal flux in physical time (as measured by some specified clock) or of temporal flux in a particular subject's subjective time.

Someone might object that this definition is inadequate because it does not fully capture the intuitive notion of temporal flux. With most of this objection I heartily agree; I concede that the definition does not capture every feature of time that someone might consider an aspect of the flow of time. Certainly it does not embody the entire content of Bergson's idea of duration. (Some of the content of that idea may be inexpressible in ordinary language.) However, the claim that my definition is inadequate on these grounds misses the point. That definition was intended only to capture the experienced *transiency* or *ephemerality* of time — the fact that moments have the property of giving way to other moments. Adopting a term which other authors have used somewhat differently, I will call this latter feature of time its *transitionality*.¹³

Understood in this way, temporal flux is not reducible to relationships of temporal order or distance among instants, dates, or events. The before-and-after relations among events are insufficient to guarantee the existence of flux.

Temporal flux, or time's transitionality, consists in the fact that each moment of history is, by nature, a transitory entity — an entity which, by its nature, must give way to new moments. It does not consist in the fact that different times or events stand in some ordering relation to each other.

Temporal Flux in Sense Experience

My aim in this section is to show that temporal flux, as I defined it in the last section, is a feature of the world as it seems. More precisely, I will show that the apparent physical objects which exist for our awareness are in temporal flux with respect to our streams of subjective time (if these objects exist at all). This conclusion is supported by certain conspicuous facts about our everyday sense experience. It also is supported by current theories of physics, which describe or summarize the behavior of the apparent physical world.

Consider what happens when you are standing on a street corner and you see a rapidly moving car. If you are like most people, you notice that the car is moving. Motion is change in an object's position over time, so you could discover that the car is moving by recording the car's positions at different times and drawing a conclusion from the data thus obtained. But you do not have to do this — at least not consciously. You simply notice that the car is moving.¹⁴ It *seems* to you that the car is moving. In other words, it is the case for some of your consciousness events

that motion is occurring.

Although humans can perceive motion, this perception may depend upon the neural integration of many stimuli. Some animals apparently have specific kinds of visual sensations of motion as well. Some vertebrate nervous systems — those of frogs, for example — appear to have a capacity for seeing motion which is not shared by humans, at least not in the same degree. Frogs, for example, can see moving bodies well; their eyes (not only their brains) can react specifically to variations in brightness.¹⁵ In view of the differences in reasoning capacity between humans and frogs, it is implausible to attribute frogs' motion perceptions to conscious inference of any kind.

To notice that an object is moving, one does not look at that object at successive moments and use descriptions of the resulting observations to deduce that the object is moving. One simply can see *now* that the object is moving. One does not have to think consciously about the fact that the object's position now differs from the object's position later. This fact suggests (though it does not strictly imply) that the movingness of a sensible object is a property of the state the object is in *now* — that is, at a single time. It is not merely a property of a set of consecutive positions in the object's history.¹⁶

If the time with which we are concerned is subjective time as described in Chapter 5, then this suggestion is correct. The fact that an object is moving can be the case for a single consciousness event. The motion of an object in the visual field does not consist simply in the occupation by that object

of different positions in the field at different times. Motion sometimes is part of the way things seem.

(A physicist might object to the preceding paragraph on grounds of relativity: since motion is relative, the objection goes, there is no such thing as a property of "movingness" which an object can seem to have. This objection disappears when we note that the required property of "movingness" is a perceptible quality and hence is relative to a subject. Since the subject has a location and a velocity, the motion relative to the subject is indeed a relative motion.)

The above argument that sense objects are in temporal flux raises some questions. The argument shows that physical objects are in temporal flux in *subjective* time. But this does not show that physical objects are in temporal flux in *physical* time — that is, in time as measured by clocks. (Such time can be measured even if one is not prepared to assume the objective existence of physical objects. One simply uses the apparent clocks in one's experience.) This kind of temporal flux is harder to establish using ordinary experience, because it is hard to ascertain, from ordinary experiences, whether an object is in the kind of state required by the definition of temporal flux. An object cannot be in temporal flux unless it is in a certain sort of state at a single time, but one cannot observe an object for just one instant of physical time. There are physical as well as physiological reasons why this is impossible; events in the nervous system require finite amounts of time, and the uncertainty principle seems to entail that a physical process involving finite amounts of energy cannot be localized

precisely in time.¹⁷ (The uncertainty principle allows much more precise time measurements than do the limitations of the human nervous system; recall what I said in Chapter 3 about the time required by a consciousness event.) In looking at a moving object, one sees the object as it goes through more than one instantaneous state. Hence one cannot be sure that anything that one sees can be taken to be a property of the object at *one* such state, rather than a property of a short segment of object history.

Some facts about the physics of moving objects suggest that motion is a property which an object possesses all at once, at each single instant of time. If one could look carefully enough at a moving car, one would find features of the car which could yield information about the car's state of motion at an instant. For example, a spinning wheel experiences internal stresses due to effective centrifugal forces engendered by its rotation.¹⁸ These stresses distort the shape of the wheel. A complete list of all the physical properties of a moving car at a fixed time would have to include the stresses in its wheels and the resulting distortions as well. Even if time were composed of true instants of zero temporal length, and even if one could take a snapshot of a moving car at a single such instant, one still could notice this distortion in the snapshot. There are other instantaneous features of this same general sort — for example, the distortion of the car's tires due to their rolling friction on the road.¹⁹ If one digs deeply enough into physical theory one can find even deeper features of this kind. The relativistic

contraction of an object's length is one such example; another is the object's quantum mechanical de Broglie wavelength. (For normal macroscopic objects this is too small to measure by ordinary means, and for objects moving at ordinary speeds the relativistic contraction due to motion is well known to be practically immeasurable.²⁰)

The view that an object's state of motion is part of that object's instantaneous state is deeply embedded in physical theory. Classical physics makes extensive use of the concepts of velocity and momentum, which are properties of the object's instantaneous state of motion.²¹ At a given time in its history, a classical particle has such-and-such a value of velocity and such-and-such a value of momentum. In relativity theory, the concepts of 4-velocity and 4-momentum supersede the classical concepts of velocity and momentum; these "4-vector" quantities, unlike their classical counterparts, are the same in all frames of reference.²² According to quantum mechanics every physical object possesses a de Broglie wavelength. This wavelength is a measure of the spatial variation of the object's wave function, which characterizes the object's state at a fixed time. In quantum theory, an object's de Broglie wavelength is inversely proportional to the object's momentum. Thus the object's state of motion is fixed by a property that the object can possess all at once at a single instant.

These examples from physics are not really necessary to my argument. They are meant to illustrate the claim that a moving physical object has, at each instant in its history,

properties which no stationary object can possess. Even without the assumption that Newtonian, relativistic, or quantum mechanics is correct, one can find in everyday experience examples of properties possessed at an instant by moving objects and not by stationary objects. The distortion of tires is among the most mundane of these examples.

Thus, ordinary experience and physical theory both agree that the state of motion of a physical object at a time is a property which the object possesses at that time. This property is an aspect of the object's state at time t . This finding has important consequences for our understanding of the nature of motion. States in the career of an object which moves relative to some observer are not the same as any states which might be found in the careers of objects which do not move. In other words, the history of a moving object cannot be built up from states taken from the histories of nonmoving objects (or of objects in different states of motion). The state of an object at a single time includes the object's state of motion. This state of motion is not merely a property of a stretch of the object's history during which the object actually moves. Instead, it is a property which the object has in its entirety *now*. One might say that a state of a moving object now is a *moving-object state* which is not identical to any *resting-object state*. The career of a moving object (moving in a given reference frame) cannot be divided up into a series of states indistinguishable from states of objects at rest.

We have arrived at one of the key ideas of Bergson's view of time, or perhaps at an updated version of that idea.

Bergson recognized that a motion is not simply a sum of instantaneous states of rest.²³

The state of motion of a physical object is a property whose possession by that object implies that the object will enter states other than its present one.²⁴ (Of course, it will not enter those states if the object loses that property. An object can lose its state of motion by being broken to bits, by being stopped by a wall, or in other ways.) Any object that has a property of this kind is in temporal flux as defined in the last section. Therefore, every moving object is in temporal flux. But every sensible object of which we know is in motion in some frame of reference or other; ordinary experience tells us this (you always can start moving away from a stationary object), and physical theory concurs. Hence every sensible object of which we know *endures*, in the Bergson-reminiscent sense of being in temporal flux.

How Events Happen

We all know that events happen. If there is anything safe that we can say about events, it is that they happen. But the fact that events happen is not as simple as it seems. It is a fact which we could understand better than we now do. To begin to understand it better, we must first understand what it means for something to happen.

When we say that events happen, we are asserting that events have a particular property. This property stands in a close logical relationship to the property of temporal flux.

The gunshot example in Chapter 4 gives us a hint of the nature of this property. First the awareness of the noise happens, then that awareness fades away as other things (including continuance) begin to happen. The earlier consciousness event is over and done with, and *gives way* to a later consciousness event. In short, the first event doesn't remain present statically, like a stone statue — it happens and is replaced by *something else*. This transitoriness of a perception of a loud noise is closely analogous to the transitoriness of object-states in the history of an object. If C and D are consecutive parts of an observed event, then one can notice, not only C and D, and not only the fact that C precedes D, but the turning of C into D, the giving way of C to make room, as it were, for D.

This fact about observed events lies at the bottom of the truism that *events happen*. When we say that an event happens, we are saying, more or less, that its earlier stages give way to its later stages. Just as a persistent object engages in *enduring* when earlier moments in its history yield to later ones, so an event engages in *happening* when its earlier temporal phases yield (in a similar way) to later ones. It is plausible to think of happening as the defining property of the class of events; events are just entities that happen.

We can define happening by analogy with our earlier definition of temporal flux.

x happens if and only if *x* has a part *y* which is present at a time *t*, and is such that the fact that *y* is present at *t*

necessarily implies that some other part z of x will be present at some time t' later than t .

Again, we can define this with respect to either physical or subjective times. The necessary implication is of the same sort used in our earlier definition of temporal flux.

It is important to note that both temporal flux and happening are defined in terms of temporal notions — namely, existence at a time and the relation of temporal order. Happening and temporal flux may not be reducible to the temporal ordering of events, but they do logically require the existence of that ordering. Thus we cannot correctly speak of something happening or enduring without specifying a series of times in which it happens or through which it endures. An entity might happen or endure in some subject's subjective time. It also might happen or endure in objective time, with the times in the definition of happening and duration read as times on some clock. In the sequel I will be concerned mainly with happening and temporal flux in subjective time. In this case, the "moments" or "instants" of time will be consciousness events. Sometimes I may make general statements about happening and duration without specifying a particular set of times. In these instances, the statements are intended to be applicable either to subjective time or to clock time.

Some Technical Notes on Instantaneous Events

The above definition of happening has two important technical consequences for the metaphysics of events. One is that instantaneous events do not happen, though they may exist. Another is that the happening of an event is not reducible to relationships of temporal order among the temporal parts of that event. I will take up these consequences in turn.

According to our definition of happening, the happening of an entity requires the possession by that entity of at least two parts which exist at different times. This rules out the possibility that an instantaneous event happens, if by an "instantaneous event" we mean an event which is localized entirely at a single instant in some series of times. (More precisely, it rules out the possibility that an event happens with respect to any series of times in which that event is instantaneous.) Perhaps the above definition of happening arbitrarily leaves out this peripheral case of instantaneous events. If it does, then my earlier statement that "events are just entities that happen" must be amended to read "events are either entities that happen or instantaneous events." (Alternatively, one could say that instantaneous events, if there are any, are not really events.) But the definition of happening does not beg the question of the existence of instantaneous events, if by "instantaneous event" one means something like "entity occupying only one instant of time."

All that is at stake is whether such instantaneous entities can correctly be said to happen — not whether they can be.

The view that a genuinely instantaneous event does not really "happen" seems reasonable. An event which truly takes no time does not involve any change or any duration; nothing goes on while such an event is present. An instantaneous event simply *exists* briefly; nothing actually *happens* during its presence. It does not even have a normal beginning or ending. The moment of its "ending" is that of its "beginning"; its ending and beginning are not distinct phenomena, and are coextensive with its presence. It is less confusing to regard such an entity as a very short-lived entity which does not embody any happening, than to regard it as an event of a pathological sort.

All this does not affect in any way the use of instantaneous events as a convenience in theoretical physics. In common physical usage, an event of zero duration represents an idealized limiting case of a very short-lasting and spatially small event. A real event need not be assigned a precise, single location in spacetime, so long as events can be located as closely as is necessary in practice.²⁵

Another consequence of my definition of happening is that the happening of an event does not consist simply in the event's being extended in time — that is, the possession by that event of earlier and later temporal parts. The temporal parts must be there, but they also must yield to each other in a particular way or the event does not happen. Happening, as defined here, does not consist in a particular arrangement in time of the parts of an event. It requires the existence of

temporal parts in an event, but is not reducible to temporal order. The temporal parts in the definition of happening are transitional or ephemeral, in the same way that the object states in the definition of temporal flux are transitional or ephemeral. Earlier I said that temporal flux is a property that an object has now, all at once; this property is not reducible to a mere temporal ordering of object states. A similar remark is applicable to happening and events. Happening is a property that an event has now, all at once, in the sense that a very brief present temporal part of the event can have the transitionality that confers happening on the event. Transitionality is not reducible to a mere ordering of temporal parts.

Incidentally, there is no guarantee that all the temporal parts of an event happen. In principle, there could be non-happening parts of events. Every event would have such parts if time were continuous and every event were a continuum of instantaneous events.

More Technical Notes: Duration and Happening

Earlier I said that happening has close logical ties with temporal flux. Now I will explore some of those connections.

An entity in temporal flux at time t is in a certain state at t . The fact that it is in that state implies that later the object will be in some other state, if it exists at all after t . A

happening entity has a temporal part, the existence of which implies the existence of some future temporal part. Thus the concepts of happening and of temporal flux are somewhat symmetrical to each other. This is not surprising, since both concepts capture the same underlying feature of reality — the transitionality of time. That things are in flux is equivalent to the transitionality of stages in the histories of things; that events happen is equivalent to the transitionality of events' temporal parts.

It is not hard to show that if there is a happening event which is part of an object's history, then that object is in temporal flux. The following argument shows this. Suppose that persisting object *O* has an event *E* in its history, and that *E* belongs entirely to *O*'s history (that is, every temporal part of *E* is part of *O*'s history). At very least, *E*'s being in *O*'s history implies that while *E* is happening *O* has a certain property — that of having an event of such-and-such a kind happen to it, or of being involved in an event of such-and-such a kind. Each of *E*'s temporal parts E_1, E_2, \dots confers analogous properties upon *O* at various times. But *E* happens; hence the presence of E_1 necessarily implies the later presence of another temporal part of *E*, say E_n , which is not identical to *E*. This implies that one of *O*'s states includes a property (being involved in E_1) whose possession by *O* entails that a later stage in *O*'s history (when it is involved in some E_n later than E_1), *O* will be in a state different from *S*. Thus *O* is in temporal flux.

The converse conclusion — that if an object is in

temporal flux then it has happening events in its history — follows if one assumes that an object in temporal flux passes through temporally extended object-*stages* (as well as object states). If an object-stage is temporally extended, then it has shorter object-stages as temporal parts; whether or not these shorter stages happen, they do give way to each other in the manner described in the definition of happening. It follows that the career of an object in temporal flux, if it contains any object-stages at all, contains temporal parts that happen. Mellor, considering objects which do not change, has argued that "thing-stages" (what I have been calling object-stages) should not be regarded as events.²⁶ But even if one does not regard object-stages as events, one can derive the weaker conclusion that an object-stage *happens* if it is in the career of an object in temporal flux.

If object-stages happen but one chooses not to count them as events, then one could further revise my earlier remark that "events are just entities that happen," to read "events are either entities that happen or instantaneous events, and entities that happen are either events or object stages."

Tense, Temporal Motion, and Time's Flow

In this section I will point out some differences between temporal flux and two other temporal notions to which it is conceptually close. These are the notions of *tense* and of the *movement of time*. (I already have discussed the movement of time, though briefly.)

The problem of the reality of tense is one of the central problems in the philosophy of time. I will not attempt here to discuss this problem in detail. Instead I will refer the reader to the literature, and in particular to the works of John McTaggart and of Mellor, for discussions of the problem and of proposed solutions.²⁷ Occasionally I will use the terminology of "A series" and "B series" which these authors have used. First I will explain these terms as I understand them.

The A series is the series of tenses, or possible positions relative to the present, which an event can occupy. The main tenses are the past, present, and future tenses; hence the past, present, and future are parts of the A series. So are more specific tenses, such as yesterday and five days ago. Since both yesterday and five days ago represent the positions of certain events relative to the present, they both are tenses in the philosophical sense. They form parts, as it were, of the past tense.

The B series is the series of all *dates*, or absolute positions in time, which events can occupy. The year 1900 is an example of a date. There is nothing that makes a particular date intrinsically past, present or future. During 1900, the year 1900 was present; in 1901, that year was past, and in 1899 it was future. Other examples of dates are: exactly midnight on January 1, 2001; the day the *Titanic* sank; and the microsecond during which a particular nuclear reaction occurred. All of these are dates rather than tenses; they can be described without reference to the present.

If the flow of time is objectively real, then there is more

to time than just the temporal ordering of events. The apparent existence of temporal flux (in the naive sense, not my rigorized one) suggests that events really do pass from being future to being present and from being present to being past. This lends great plausibility to the commonly held view that temporal flux depends upon changes in tense.²⁸ The transiency of moments, which we call time's flow, seems to involve in an essential way changes in the tenses of dates and of events. Without changes in tense, it seems, there could be no real passing of the present into the past. Thus the existence of real flux might be thought to require the existence of an *A* series of real tenses.

Despite its plausibility and popularity, this latter belief is incorrect. *Neither temporal flux nor happening requires the reality of tense.* To see why this is so, consider the history of a persisting object in a hypothetical tenseless universe. This history is composed of a series of object states. There is no reason why one of these states could not possess some feature which compels other states in the object's history to exist. For example, the object state *S* might be such that an object which is in *S* has a certain quantity of energy. Suppose that the law of conservation of energy holds in the world we are imagining. Then if an object is in object state *S*, there must be an object state of some kind after *S*, for otherwise the energy of the object in *S* would be lost at times later than *S*. Even if the object disintegrated, there would have to be something afterwards which possessed the "lost" energy. And as long as the object itself continues to exist, its being in a given momentary state necessarily

implies that there is a later state; this later state cannot be identical to the earlier state due to its temporal relationship to that earlier state. Thus, an object in a tenseless universe could be in temporal flux without undergoing any changes in tense. An event in the history of such an object could be a genuine *happening*, could be genuinely ephemeral and transitional, without undergoing any changes in tense.

The transitionality of the moment embodied in the concepts of temporal flux and happening does not depend upon the existence of changes in tense. Temporal flux and happening are properties that an object possesses *regardless of its real or apparent tense*. The essence of temporal flux and of happening lies not in some present entity's becoming past, but in the fact that some entity compels another entity to exist in the future. Both entities really embody transitionality, happening, or flux — but these entities do not need to change in tense, or even to have tenses, to do so.

The existence of flux and of happening does not require any actual movement of time, such as a movement of events from the future toward the past. Any such actual movement requires the existence of tense in some form. Since flux and happening do not require tenses, they do not require this movement.

The reality of flux and happening is equally compatible with the reality of tense and with the unreality of tense. The existence of entities (things or events) that are in flux, or of entities that happen, does not require or exclude the possession by those entities of tenses, or the movement of those things from one tense to another. Hence the thesis that

flux and happening are real cannot be used to argue that tense is or is not real, and neither the reality nor the unreality of tense has much bearing on the reality of flux and happening.

Of course, if the existence of tense is contradictory, as McTaggart and Mellor have claimed,²⁹ then *nothing* is logically compatible with the existence of tense. The arguments of McTaggart and Mellor against the reality of tense show that the concept of tense raises certain logical problems. One does not have to agree fully with the conclusions of these arguments to recognize this problem.

The connection between tense and the reality of the past and future bears mentioning in this context. According to many people's feelings, the past does not exist anymore and the future has not yet come to be. Many people feel that neither the past nor the future is fully and genuinely real.³⁰ This outlook amounts to a rather overwhelming version of the thesis of the reality of tense; clearly it implies at very least that there is a present tense. The existence of temporal flux and happening does not depend upon the fate of this outlook. The definitions of temporal flux and of happening do not require us to hold that future or past objects exist; whenever they speak of future or past entities, they can be revised to speak of entities that will exist or have existed, rather than to future or past entities which do exist. One can apply these definitions to things and events regardless of whether the future or the past is real. Also, the existence of temporal flux does not entail that past and future events are in any sense less real than present ones. The fact that an

object is in flux does not make any of that object's states or stages less real than any others.

Summing up some of the conclusions of the last few paragraphs, we can say that the reality of temporal flux is independent of the reality of tense and of the movement of time. Flux, happening, and transitionality can exist in a world in which only the *B* series, and not the *A* series, is real.

In a world in which there is no *A* series, every feature of time can be reduced to something which exists tenselessly and unchangingly. A real flow of time closely resembling Bergsonian duration is precisely the sort of feature one would least expect to find in such a world. One might object on these grounds to my thesis that transitionality can exist in a tenseless world. However, this objection misses the point, for reasons set forth in the last few paragraphs.

One easily can turn this objection around to show that no *B* series theory of time can be correct *unless* there is transitionality in the world. Happening is not a byproduct of the temporal ordering among events; it is a feature of the individual events themselves. An event could happen even if it existed in a universe lacking any other stages or events besides its own temporal parts. Every *B* series approach to time of which I am aware takes the notion of event, or some related notion like time or date, as fundamental. On some views, instants are constructed with the help of events³¹; at very least, times are identified and located with the help of events. No *B* series theory (at least none of which I know) categorically denies the reality of all events except non-

happening instantaneous ones. Thus, *B* series theories of time are theories about a world populated by temporally extended events — in other words, by *happening entities*. Therefore, any attempt to reduce away happening and related "fluid" features of time with the help of a *B* series account of time will fall into a vicious circle. As soon as one uses the notion of event in the foundations of the theory, one imports happening into the theory too — for an event (with the one possible harmless exception discussed earlier) is simply an entity that happens.

The reality of temporal flux or of happening does not imply the reality of tense, with all the possible logical headaches which that entails.³² Nor does it imply any sort of "motion" of the present into the past or future. It does not imply that there is a "flow of time" in any sense requiring the existence of an *A* series. The flow of time is due to the happening of events and to the temporal flux of things, not to any relationship between objects and tenses.

Re-Charting the River of Time

To close this part of the discussion of time's flow, I will present a metaphor illustrating the difference between the tensed view of time and the view of time's flow for which I have argued.

Consider the well-known picture of the "river of time."³³ On the tensed view, there is a real flow of time; things floating on the river are carried downstream by the current.

An event is much like a fallen leaf on the river's surface, which is first upstream, now here, later downstream. In this manner events emerge from the future, reach the present, and disappear into the past.

The above metaphor embodies the *A* series view of time. On standard *B* series views, the river is replaced by a perfectly still pond. In this pond things just float on the surface without going anywhere. There are directions called "upstream" and "downstream" in the pond, but these directions are demarcated by something else besides a real flow. (Perhaps there are strings of clocks on the pond, with the downstream clocks in a string reading later times than the upstream ones. Perhaps the strange creatures in the pond simply are afflicted with *feelings* that the pond is a river.)

On my view of time, time is more like a river than a pond. The flow of the river, with all its restlessness and energy, is quite real. But events are not boats or autumn leaves continually floating downstream. Instead, these entities are like ripples in the water surface created by underwater rocks. Such ripples can be quite persistent, as anyone who has watched shallow streams knows. The ripples do not move from upstream to downstream, but the river really flows nevertheless. The ripples themselves, though stationary with respect to the river's banks, are filled with a restless flow. In this sense, they truly are moving, although their positions along the river never change.

The Objective Reality of Time's Flow

In the preceding sections I have argued that temporal flux and happening are properties of physical things and events. My conclusions there suffice to establish the reality of temporal flux and happening, granted the reality of the physical world. In this section I will explore another avenue to the conclusion that temporal flux and happening are real. I will argue that we can establish the reality of temporal flux and happening in subjective time without assuming the existence of either physical objects or physical events.³⁴

To notice that an object is in temporal flux, one must notice that the present state of that object is giving way to new object states in the manner which I described in the preceding sections. To notice this, one must at least notice the state into which the object is passing. The perception of temporal flux in an object involves perception of the yielding of one state to another. This requires noticing that a state is "going away" and is being replaced by a new state. An experience of glancing at an object in a fixed state, without seeing it starting to enter into new states, is not an experience of flux. To see an object as being in flux is to see the object as it begins to move into new states. (One could try to create a counterexample to this by staring at an object which remains qualitatively unchanged, but even that experience is one of an object in flux — one in which states are being succeeded by new states that happen to be

qualitatively similar to each other, though standing in different temporal relations to those earlier states.)

Similar remarks hold for temporally extended events. To watch an event is to watch an event *happen*. To watch a fire is to watch a fire burn, at least for a moment. One cannot notice an event without seeing something happen, that is, without seeing a transition occur between brief temporal parts of the event — without seeing one brief temporal part transformed into another.

These facts about experiences of temporal flux and happening are sufficient to allay all doubts about the reality of these two properties, at least as they are found in *subjective time*. Consider the following description of what happens when a subject perceives that a temporally extended event happens. During a particular consciousness event E, a temporal part y of the event is perceived; in the consciousness event F which succeeds E in the subjective time order, another temporal part z of the event is perceived. The perception of the event must involve the perception of z's beginning to be present and of y's ceasing to be present; without such appearances, it would not seem that the event as a whole happens. Thus, the perception of y cannot occur *in full* — cannot be *finished* — without the perception of z. The fact that y is perceived implies that z is perceived in some manner or to some degree. Hence the subject's experience during the total event has two temporal parts in subjective time, and the existence of one part of the experience requires the subjectively later existence of the other part. Therefore, experience *happens* in the subject's

subjective time.

The upshot of all this is that an experience which seems to be a perception of a happening entity itself happens in subjective time. If that experience really is a perception of an event, then that event really happens too, at least in subjective time. This argument can be extended to cover temporal flux as well as happening (just replace perception of events with perception of objects which are in certain states).

G.J. Whitrow has argued for the claim that our experience of the passage of events cannot occur without a real passage of time.³⁵ Thus, my conclusion is a variant of Whitrow's, although Whitrow's argument (and conception of the passage of time) are rather different from mine.

My argument for the reality of temporal flux and happening in subjective time bears strong similarities to an argument which Mellor once deployed for a different purpose (to analyze our experience of the order of time).³⁶ It differs from Mellor's argument in that it emphasizes causality much less and links the objective temporal order of experiences much less directly to their apparent order. The argument presented here shows that an appearance of duration or temporal flow is an appearance of a special kind, having a metaphysical significance greater than that of appearances of any other sort. One can have a visual hallucination of an elephant even if there is no elephant. But to have an illusion that things are enduring or that time is passing, one must perceive falsely that one moment is passing into another, that objects really are moving, that an

event is transpiring, or the like. One would have to have an illusion of successive events or stages — an illusion which itself consists of experiences which are successive (at least in subjective time), an illusion which consists of a transitional succession of illusions. In brief, to hallucinate flux one must have a *flowing hallucination*, and to hallucinate happening one must have an experience that happens. Once I have hallucinated that an event has happened, already something has happened; my experience of the event happened. A similar remark holds for the hallucination that time has passed (another variant of Mellor's argument³⁷ could be used to establish this). The illusion of time is not an ordinary illusion because *time is implicit in the very act of hallucinating time*. Hence time — or more precisely, the phenomena of temporal flux and happening — cannot be entirely unreal.

Thus the apparent flow of time cannot be entirely a byproduct of structural features of a fluxless, happening-free world. Time's flow, as we perceive it, cannot emerge from an entirely fluxless substratum. The transitionality of time is an objective feature of the world. This does not imply that time has an extramental origin; I have not ruled out the possibility that flux or happening is a product of mental phenomena. (Note also that my argument for the reality of transitionality showed that flux and happening are real in *subjective* time. The argument also will work for clock time, if the clock goes forward in the usual manner while subjective time progresses.) But even if these temporal features of the world were psychological in character, *they*

still would be objectively real.

We have taken an important step along the path from experience to world. We have shown, using facts about how things seem, that *experiences happen*. We also have shown that events which exist for observers happen (if those events are real), and that objects which exist for observers are in temporal flux (if those objects really exist). In short, we have passed from how things seem to the *reality of time's flow*. Along the way, we have shed some light on the real nature of that flow.

Besides establishing the reality of time's flow, the above argument shows that time has a unique epistemic status among all the general features of the world. It shows that one cannot apply the method of Cartesian doubt to time as one might apply it to material objects. Attempts to doubt the reality of time as Descartes doubted that of the material world can lead a skeptic only to the certainty that time — in some form, either subjective or objective — is real.

From Consciousness to Objective Time

The conclusion that experiences exemplify happening has two important consequences.

First, it implies that the presence of real happening in the world is a necessary consequence of the existence of subject histories of a certain sort. If there is a subject history which has more than one consciousness event, and a temporally extended event exists for a string of consecutive

consciousness events in that history, then there really are events which happen, at least along the subjective time coordinates of observers. In short, *a world containing consciousness like ours is a world in which there is time of some sort.* There is no need to look further than consciousness to explain why time is a feature of reality.

Secondly, we now know that any conscious being which has more than one consciousness event in its history *really lives through time of some kind.* This is the subjective time dimension whose moments or "points" are the subject's consciousness events. Subjective time is characterized by a kind of transitionality, and therefore really is a variety of time.

A world containing experiences of the sorts with which we are familiar has to be a world in which time is objectively real. The reality of such experiences is a sufficient condition for the reality of happening, and therefore for the reality of time in its broadest sense. (This does not mean that the existence of consciousness implies the existence of time as physics knows it, with all of its ordering and metrical properties. It means only that the existence of a consciousness event with a successor is sufficient for the existence of bare happening, which is the most central feature of time.)

These results also serve to clarify the logical connections between mental concepts and temporal concepts. Although philosophers long have noticed that there are such connections, the nature and extent of those ties remains largely mysterious.³⁸ In this chapter I argued that any real

objects which we perceive as happening or in flux in subjective time really are happening or in flux. Thus, objective temporality is a facet of any world in which conscious observation occurs. This connection between consciousness and objective temporality is a consequence of the nature of the peculiar mental circumstances which give rise to subject histories. Roughly speaking, a prolonged event of observing or experiencing a persisting subject must have parts which "give way" to other parts in a way that constitutes happening.

This kinship between consciousness and temporal flux implies that any conscious being has objective temporal properties. Such a being can be said to exist in time, provided we do not mean by this a form of time closely resembling the physical time with which we are familiar. The experiences in the history of a conscious subject really happen. Even if this happening is only happening in subjective time, it still is objectively real. If the subject's subjective time runs in the same direction as physical time (recall Chapter 5), then there is happening in physical time as well.

Thus, there is no such thing as a conscious being free from all flux, duration, becoming, or happening. It is the nature of a conscious being to live in some form of time, just as it is the nature of a garden snail to live in a shell.

From Brain to Cosmos

Chapter 8

The Experience of Time

In the last chapter I discussed the problem of the nature of time's flow and touched upon the problem of the reality of tense.¹ In this chapter I will explore some logical connections between the concept of consciousness and the concept of tense. The results of this exploration will shed some light upon the more general issue of the connections between mental concepts and temporal ones. Some people regard such connections as forming a significant puzzle in the philosophy of time.²

The tensed view of time is what I previously called the *A* series view — that is, the doctrine that events really do have tenses, such as past, present, and future, instead of merely seeming to us to have tenses. It is closely allied to the commonsense view that time is moving toward the future. Some philosophers have regarded the tensed view of time as logically untenable.³ My main contention in this chapter is that the tensed view of time is consistent and true when understood as a thesis about *subjective* time. More

precisely, the main points of the tensed view come out true and consistent if one assumes that tenses are properties defined *relative to consciousness events*, rather than being absolute or objective properties of events. I will show that the picture of time typical of *A* series theories is right in many of its points, provided that we treat tense as a subjective property in this way. For example, relative to a single consciousness event, future events do not exist while present and immediately past ones do. Overall, I aim to show that the tensed view of time is a correct way of thinking about subjective time, even if it is false when applied to objective time. Tense is a real feature of subjective time.

The material in this chapter consists in part of technical details; because of the nature of the arguments I will use, there is no way to avoid this. The details of the arguments will be of interest mainly to those interested in the philosophy of time. However, the general conclusions reached here will be used occasionally later on. Those who wish to skip the details of the arguments can do so, though this may, of course, make the reasons for the conclusions less clear.

Subjective Tenses

In the last chapter I mentioned *A* and *B* series theories of time. According to *A* series theories, tense is real; it is a feature which times or events objectively possess.

According to *B* series theories, tense is not real, although it may appear to us to be real and may, in a certain manner, even have practical importance.⁴ *A* series theorists typically hold that events really change from being future to being present to being past. *B* series theorists typically attempt to reduce the so-called movement of time to the ordering and other properties of the *B* series of times or events. The *B* series viewpoint entails that all events, whether they are labeled past, present, or future, are tenselessly real.

One limitation of *B* series theories is their difficulty in giving a proper account of the ephemerality of events. In Chapter 7 I showed how to overcome this limitation. To some people, the tenseless view of time seems to make it difficult to justify our normal emotional feelings about the future and the past. In Chapter 7 I alluded to this difficulty; here I will discuss it in greater detail, and will show how it can be overcome.

The hypothesis of the reality of tense is beset by logical difficulties. Mellor has argued (successfully, I think) that these difficulties rule out the objective reality of tenses.⁵ Yet as Mellor has noted, tense plays an important role in human psychological life. Everyone knows that we experience the passage of future events into presence, and that we live through the seemingly irreversible slipping away of the present into the past. Mellor's explanation for the psychological role of time's apparent passage depends on his "token-reflexive" theory of the semantics of tensed beliefs. On his account, tensed beliefs may be true (and cannot be replaced entirely by tenseless beliefs) even though

there are no objective tenses.⁶ Mellor has noted the important fact that tensed belief is a practical necessity for human action.⁷

Within the subjective time order on a subject history, any consciousness event is either past, present, or future to any other. At any subjective moment in the history of a subject besides the first and last consciousness events (if these exist), there is a series of past consciousness events of that subject stretching out into the subjective past, and there is a series of future consciousness events of that subject stretching out into the subjective future. One can use these consciousness events to define a series of tenses for the subject's history, just as an *A* series theorist uses past, present, and future events to define the *A* series.

Consider a single consciousness event, *E*. If a fact is the case for some consciousness event in the subjective past of *E*, we will say that that fact is *subjectively past* at *E*. If a fact is the case for some consciousness event that is in the subjective future of *E*, we will say that that fact is *subjectively future* at *E*. If a fact is the case for *E*, we will say that it is *subjectively present* at *E*. In this way, we define three subjective tenses: the subjective past, present, and future.

We may also speak of a thing or event which is not a consciousness event as being subjectively past, present, or future for a consciousness event *x*, if the fact that that entity exists is subjectively past, present, or future for *x*.

Tense logicians know that phrases indicating tense, like

"is past" and "is future," can be treated as operators on sentences.⁸ The above definitions could be used in a straightforward way to define operators for subjective pastness, presentness, and futurity, analogous to the usual tense logical operators for the past and future tenses.

These *subjective tenses* capture certain everyday intuitions about the tenses of experienced events. They reflect a subject's living experience of time. It is a fact of everyday experience that some things, situations, and events seem to be present and that others once seemed to be present; we believe that still others will seem to be present. One really does live through the passage of one's experiences from presentness to pastness. The above definitions of the subjective tenses describe the past, present, and future as we feel them in our experience.

Note that the subjective tenses as defined here are not just *apparent* properties of facts, events, and things. They are not merely properties that these items *seem* to have. There are facts which are subjectively past to you now, but are such that it does not seem to be the case that those facts once were true. (These include facts about events which you really lived through, but which you no longer remember.) There also are facts which seem to you to have been the case, but which are not subjectively past for you now. (These include imperfectly recalled memories.) Analogous remarks can be made for the subjective future tenses. The subjective tenses reflect the real flow of experience which occurs as consciousness passes from one subjective moment to another. They need not reflect the way things seem now.

We could, if we wanted to, define other tenses corresponding to the way things seem now. We could call these *apparent* tenses.⁹ A fact is *apparently past* at a consciousness event if and only if for that event, it is true that that fact once was the case. Similarly, a fact is *apparently future* at a consciousness event if and only if for that consciousness event, it is true that the fact will be the case, and a fact is *apparently present* at a consciousness event if and only if for that consciousness event, the fact is the case. The apparent tenses can be thought of as tenses which are apparent *in the subjective world* — features of the realm of appearance of a consciousness event. The subjective tenses, in contrast, are features of facts and things in the subjective world, but are not necessarily present in the subjective world themselves. They need not always be part of the way things seem now.

For the remainder of this chapter I will concentrate on *subjective* rather than *apparent* tenses. The subjective tenses reflect the real temporal order of subjective life. Some of what will be said carries over to apparent tenses as well.

Mellor's and McTaggart's critiques of tense suggest that no event is objectively past, present, or future.¹⁰ But their arguments cannot be used to derive a contradiction from the *subjective* reality of tenses. If one replaces the tensed sentences in Mellor's argument against tense with the corresponding subjective tenses relative to some consciousness event, one finds that the revised argument will not go through. The relativization of tenses to different

consciousness events forestalls any contradiction.

By dividing subjective tenses more finely than we have done, one can in principle define a *subjective A series* of subjective tenses. If tenses turn out not to be real, then none of the tenses in the ordinary A series is exemplified objectively; that is, it is the case objectively that no entity has any of these tenses. However, this statement does not carry over to the subjective A series. The subjective tenses are properties which are defined in terms of moments of conscious life. They are not properties which an event can have in an absolute (non-relative) way. (Since the subjective tenses are not apparent tenses, an entity which possesses one of them at a consciousness event might not even possess it *for* that consciousness event.)

Subjective Tense and Subjective Ontology

Many people feel that past events already are over and done with, and that future events do not yet exist. This view is incorrect if there are no objective tenses, yet it is intuitively plausible. As I experience the passage of time, future events become present, and really do come into being for all practical purposes. As I act, I feel that I am creating events, making them happen, not just unearthing them from a temporarily inaccessible realm of future events that already exist. Even if I begin to believe that I do not really create events, still I am likely to suspect that causes bring their effects into being. I may believe intellectually that past,

present, and future events are equally real, yet I still may feel that new events are created as time passes.

The intuitive view that future events become present and thereby come into being is wrong if taken as a belief about the objective being (that is, the actual existence) of events. However, it is quite true if taken as a statement about the *subjective* existence of events. Given a subjective moment (that is, consciousness event) x in a subject's history, events in the subjective future of x do not exist for the subject's current consciousness event x . Events in the immediate past of x do exist for x . Events in the more distant subjective past of x do not exist for x , but they are connected to x via a chain of consciousness events of a specific sort: any such distant past event exists for a consciousness event, which exists for a consciousness event, which ... which exists for x . Such distant past events do not exist for x , but can be said *to have existed* from the standpoint of x . Hence the naive picture of events coming into being as time passes is correct, with few qualifications, if by "time" and "being" one means *subjective* time and *subjective* being.

The intuition that the present situation brings future situations into being also turns out to be at least partially accurate when applied to subjective time and subjective being. In particular, each consciousness event in a subject history stands in a relationship to its successor which is somewhat like the relation of causation. This is the relationship by which one event "gives way" to another; I discussed this relation in earlier chapters. Although this relationship may

not strictly be a causal relationship, it resembles a causal relationship in some respects.

These considerations suggest that Mellor's ideas about the psychological function of tense, even if largely right, do not give tense a sufficiently high place in mental life. Tense is not merely an indispensable feature of statements and beliefs; it also is a feature of the apparent facts, things, and events which exist for consciousness events. The possibility that there are no objectively real tenses does not change this. Tense is a paradigmatically subjective temporal notion.

The reality of tense *relative to consciousness* should put to rest any unease one might have about the grounding in a tenseless world of normal human attitudes towards the past and the future. So far as conscious life is concerned, the world is not tenseless at all.

Temporal Modalities: A Technical Note

I will close this chapter by indicating how an account of the notion of *truth at a time* might be extracted from the theory of subjective time developed above. I will sketch the beginnings of an account of truth at a time which seems to avoid certain pitfalls characteristic of tenseless accounts of this notion.

Intuitively, a statement can be true or false at a time. Even so simple a sentence as "The dog is black" can be thought of as having different truth values at different times. The subject of tense logic¹¹ is grounded on this recognition;

in it, tensed statements are analyzed into tenseless statements plus phrases indicating tense. For example, "The dog is black now" obviously may be rewritten "It is now the case that the dog is black"; "The dog will be black" may be rewritten "It will be the case that the dog is black," and "The dog was black" may be rewritten "It was the case that the dog is black." Standard tense logics use this sort of analysis of tensed statements. The alphabets of these logics contain modal operators which serve as formal equivalents of "It will be the case that" and "It was the case that".¹² These operators can be affixed to statements, tenseless or tensed, to form new statements. There is a standard semantics for tense logic.¹³ This semantics is analogous to the possible worlds semantics for other modal systems, with times or instants playing the role of possible worlds.¹⁴ For example, if P is a well-formed formula of tense logic, then given a model for the tense logic and a time t in that model, the formal equivalent of "It was the case that P" is true at t if and only if P is true at some time u in the model which is earlier than t in the time order relation of the model.¹⁵

Mellor has proposed a treatment of unary properties of things as relations between things and times.¹⁶ According to Mellor's account, subject-predicate sentences which normally are taken to express the instantiation of a property by an object, and which may be true or false at a time, can be taken to express the instantiation of a relation by an object and a time. On this view, "The cat is black" has essentially the same meaning as "The cat has B to t," where

B is a relation (of blackness at a time) and t is a time or date (the time at which the sentence is uttered).

Intuitively, there is a problem with this account: some people find it hard to believe that black is a relation whose field contains *times or dates*.¹⁷ This sense of disbelief is especially strong with regard to perceptible qualities like colors. One may find it implausible that colors are not exemplified by colored objects alone, but by colored objects together with *dates*, of all things.

The theory of subjective time which I proposed in Chapter 5 suggests a way to get a more natural account of truth at a subjective time, at least for some statements. Moments of subjective time exhibit some of the possible-world-like character which temporal logic imputes to times. They have this character because they are consciousness events, and hence can endorse statements as can possible worlds (recall Chapter 3). A sentence can be true or false at a subjective moment because a sentence can be true or false for a consciousness event. Thus there is a natural sense in which a sentence (or proposition) can be said to be true at a time, provided that one is speaking of moments of subjective time in a subject history and not of times of some other sort. The modal, worldlike character of consciousness events leaves room for an account of time-dependent truth which does not have the counterintuitive character of the *B* series-based recasting of properties as relations. On an account of this sort, "The cat is black" is true at subjective time t and false at subjective time u if and only if for t the cat is black and for u the cat is not black.

This sketch of an account of truth at a time can be extended to other sets of times besides subjective histories; in a moment I will show how this can be done. But the sketch as it stands already is helpful. It suggests how we might account for the truth at a subjective time of subject-predicate statements whose predicates refer to perceptible qualities. These are precisely the cases of truth at a time for which the replacement of properties by relations to times conflicts the most strongly with intuition.

As it stands, the account suggested here can give "The cat is black" a truth value at t only if t is a subjective moment involving experience of the cat. But this truth gap should not be surprising. For a consciousness event not involving experience of the cat, the cat does not exist. Therefore, the problem of filling this truth gap is simply the known problem of assigning truth values to statements containing nondenoting terms.¹⁸ There also are cases in which the cat exists but has no color at x — for example, if the noticing subject is blind and feels the cat. In these cases, the cat exists for x , but the cat cannot have any color at x .

We should be able to extend this account of truth at a time to physical times (that is, the times measured by clocks). Since I have not analyzed the notion of a physical time or of physical simultaneity here, I will not make this extension explicitly. Instead I will merely indicate how it might be done.

We could say that "The cat is black" (P) is true at a physical time t for a subject S if and only if S 's history contains some consciousness event x which takes place at t and is such that

P is the case for x. To complete this definition we would have to specify a simultaneity relation connecting consciousness events with physical events. If consciousness events are events that happen to brains, this is not a big difficulty. It seems reasonable to accept an analysis of perceptible-quality predications like "x is black" which makes this statement objectively true if and only if an observer under certain standard conditions would see x as black. Having done this, we can give conditions for the truth of "x is black" at a physical time t in terms of that statement's truth for observers at time t. Using the above analysis of truth for an observer (subject) at t, we can write the truth conditions for the statement at a physical time in terms of its truth for consciousness events.

The preceding account of truth at a time suggests that an object truly can have a perceptible quality at a time. There is no need to convert perceptible qualities into relations between objects and times. Subjective time, and probably physical time as well, inherit their intensional characteristics from consciousness events, which are worldlike.

From Brain to Cosmos

Chapter 9

Spacetime and Happening

In Chapter 7 I presented a new view of time's flow. I argued that the flow of time is an objectively real feature of the universe, but that it arises from an intrinsic "flowing" character of things and events instead of from a real movement of time. In this chapter I will discuss the implications of this dynamic picture of time for the scientific conception of spacetime. During this discussion, I will begin to present the new view of reality which I promised in Chapter 1. Although I will not unveil this view in its entirety until later, I will establish some of its main points here.

The special and general theories of relativity suggest that time is a dimension of the physical world, analogous to the three dimensions of space and sharing most of the properties of space. There is a widespread feeling that this "spatialized" view of time is incompatible with the belief that time's flow is real. I will begin this chapter by arguing that this feeling is quite mistaken. Mellor has argued that

relativistic physics does not compel us to adopt a thoroughly spatialized view of time¹, but here I will do more. I will use the ideas about time developed in earlier chapters to reconcile the existence of real temporal flux with the view of physical time suggested by relativity.

The Timeless Happening of Events

In earlier chapters I suggested (following McTaggart and Mellor) that there are no objective tenses. The objective unreality of tense implies, among other conclusions, that all events are equally real. Events do not begin to exist when they become present for us, nor do they cease to exist when they become past for us. They exist in a tenseless manner; they do not have to be present to exist.

Anything that fails to be a happening entity also fails to be an event. The one possible exception to this rule is an instantaneous event (recall Chapter 7). A non-instantaneous event has the property of happening. But past, present and future events all are equally real. Since all of them equally are events, all of them possess happening. In other words, every event — whether past, present, or future for us — *happens*. Every event that has happened, is happening, or will happen can be said *to happen* in a tenseless sense. Of course, not all of these events are happening now; when we say that an event is "happening now," we normally mean that it is present. But if events exist tenselessly and each event exemplifies happening, then we are forced to conclude

that all events happen, regardless of their supposed tense.

A similar conclusion holds for the temporal flux of persistent objects. Everything which persists through someone's subjective time is in flux tenselessly. The fact that someone correctly labels that thing as a "past," "present," or "future" object relative to some time does not change this.

All events happen tenselessly and all persisting things endure tenselessly. If events and things are embedded in a spacetime as physicists postulate, then all the contents of that spacetime possess duration or happening. In other words, *the spatiotemporal world as a whole perpetually endures*.² Its contents perpetually happen, perpetually flow. The fact that the world as a whole does not move from the past toward the future is irrelevant to this conclusion. If an event happened in 1900, then it is true now that tenselessly speaking, that event *is happening*. If a physical object existed in 1900, then it still is true now that that object *endures* in a timeless way — that it tenselessly *is* in temporal flux.

When speaking in this way, one must be very careful to distinguish between tenseless existence and happening on the one hand, and existence and happening in the present on the other. The possibility of confusions like this was known to McTaggart, who recognized that confusing such ideas or phrases can lead quickly to absurdities. The following example, though not McTaggart's, embodies his thinking on this topic; my subsequent claim about the source of the confusion is essentially his.³ If the last battle of World War

It happened in 1945, then it is the case now that that battle *happens*, for it is an event and events tenselessly happen regardless of what date is now present. But it goes without saying that that battle is not now a present event. The sentence "The battle is happening now" normally means that the battle happens and is present. But we also know that now (and at any other time), it is tenselessly true that the battle happens. Confusion between these two kinds of "present" happening can lead to brazen absurdities like this: the battle happened in 1945; but everything that ever happened happens tenselessly; therefore, the battle is happening now, so it is a present event. The mistake here is confusion between the idea that the battle is happening tenselessly (so that now it is the case that it is happening tenselessly) and the idea that the battle is located at the present date (so that now it is the case that it is a current event). The use of phrases like "now the battle is happening" can lead us into this confusion. We must be exceedingly careful to avoid this kind of confusion if we think that events which are not present may still be items that happen.

Since every event in spacetime tenselessly happens, it follows that the entire history of the physical universe tenselessly happens. Spacetime as a whole has the feature of happening. The following example — which is not to be taken as anything more than science fiction of the most venturesome sort — illustrates what this means. It suggests a way to think about the timeless happening of events — and a new way to think about the concept of spacetime.

Suppose that one somehow could get outside of spacetime and observe spacetime as a whole — including all times at once — from a vantage point outside of spacetime. Suppose, for instance, that there were another universe, of the sort which physicists speculate about today.⁴ Further suppose that someone in that universe somehow managed to observe our entire spacetime, or some large region therein, from the outside. If the standard physical conception of spacetime (or any reasonably close variant of it) is correct, then the external observer would see a four-dimensional world containing events. (Some current physical theories use a world with more than four dimensions,⁵ but this does not affect the argument as long as there is a dimension of time.) Ordinary motions, and the histories of objects which persist through physical time, would appear as paths (worldlines) having extension in the time direction. But aside from these motions and histories, which would appear as paths rather than as motions, our observer might notice the spatiotemporal world itself happening, with its contents giving way to further states of themselves. If the observer could detect the property of happening, then the observer could see the world, including all times, going on and on and on. One would see this because the contents of spacetime *happen*. This happening of the contents of spacetime is a temporal phenomenon in a broad sense. It is, so to speak, a persisting of spacetime itself through time of some sort. However, it is not a *motion* which can be represented by a series of states in time and summed up by a worldline in spacetime. (Of course, one cannot take this

metaphor of persistence too literally, since spacetime is not occupying successive moments of time — but more on *that* possibility later.)

This picture of a timelessly happening spatiotemporal world suggests that the conceptual gulf between the concept of happening and that of the *B* series is rather wide. The *B* series of dates, as measured by any given clock, lies within spacetime. Ultimately its dates are just positions along some time axis (or reference worldline) within spacetime. Items could have dates in that series even if there were no happening. A geometric structure similar to spacetime, but with happeningless four-dimensional bodies replacing events, would still have a fourth coordinate; hence points in it could be given dates, and a *B* series could be defined. But in the real world, in addition to the coordinate or direction in spacetime ordinarily called "time," there is *another aspect of time* which is neither a direction nor directly reducible to spacetime's geometry. This is the time in which the spatiotemporal world might be said metaphorically to move. This aspect of time is *happening*, an objective feature of the spatiotemporal world.

Global Happening and Super-time

Some philosophers have taken seriously the equation of temporal flux with a motion of spacetime. Broad once mentioned and rejected the possibility that the present moment travels forward through time as it seems to do.

Such a movement, Broad realized, would have to happen in a dimension of time other than the time dimension currently known to physics.⁶ George N. Schlesinger has called this hypothetical extra dimension of time "super-time."⁷ My conclusion about the timeless happening of spacetime does not require an extra time dimension, though it certainly does not rule out the existence of such a dimension. I will make no use of the idea of an extra dimension here. Nevertheless, the pure happening of the universe which I portrayed above might be regarded as a "non-metrical super-time" — a form of temporality within which the flow of ordinary time occurs, but which cannot be regarded as a dimension or as extended, and cannot be measured or quantitated by means of a time coordinate.

Broad rejected a super-time interpretation of time's flow because it threatened to lead to an infinite regress:⁸ if one tries to explain the flow of time as a motion which takes place in super-time, then one still must explain the flow of super-time; that explanation demands a super-super-time, and so forth. But a non-metrical notion of super-time can block this regress. I have argued that there is a sort of temporal flow which is not reducible to movement through a series of times. This flow is not associated with changes in a time coordinate; an object can participate in this flow without going through successive instants of time (or, for that matter, of super-time). To stop the regress of super-times, one need only regard one of the supposed extra time dimensions — super-time, super-super-time, and so forth —

as forms of non-metrical temporality of this sort instead of as a new quasi-spatial direction. If one regards spacetime's "motion" as non-metrical temporality in this way, then there is no need to postulate a super-super-time in which this movement occurs. The flow of non-metrical "super-time" cannot be conceived of as a movement in a super-time direction or down a super-time axis. There is no direction associated with such a "super-time." (Of course, one can end the regress at a later stage if one thinks, for some reason, that there are extra time dimensions.)

Non-metrical "super-time" of this sort is simply the timeless happening of the contents of spacetime. The view that everything in spacetime timelessly happens may be regarded as a refinement and relativization of the super-time theory. This new view postulates a kind of time beyond the time we ordinarily measure, but it does not postulate extra geometric dimensions of time as the super-time theory does. (Such extra dimensions may or may not exist, but their existence is neither necessary nor sufficient for explaining the fact of time's flow.) The view that spacetime tenselessly happens is analogous to the super-time theory in some respects but differs from it in other crucial ways. Both ideas postulate a temporal aspect of reality above and beyond the time dimension known to physics, but according to the non-metrical view, this aspect is not an extra dimension of time. Instead it is a non-geometric feature of the world — what I have called tenseless happening.

Subjective Flux Is an Aspect of Tenseless Happening

In Chapter 8 I argued that tense and changes in tense exist relative to consciousness events. Now we find that subjective changes in tense are part and parcel of the tenseless happening of the spatiotemporal world as a whole. Subjective changes in tense accompany the occurrence of experiences. The fact that experiences happen ultimately is an instance of the tenseless happening of events. Hence the perceived flow of time (change of viewpoint along the stream of consciousness) can be thought of as an aspect of the real tenseless duration of the world.⁹

Everyday experience reveals that subjective time passes as physical time passes. Both forms of time "flow" in the same direction. The above remarks on subjective time suggest that this empirical conclusion may not be logically necessary, or at least does not follow from the definitions involved. (Whether it is required by physical laws is a separate question which I will not take up here.) The subjective movement of time depends upon the tenseless happening of the consciousness events in a subject history. However, it does not require those consciousness events to be situated in any particular way in *physical* time. The experience of time's flow does not appear to depend logically upon the orientation of the subject history in spacetime. Probably it is not logically necessary that the

string of consciousness events which make up a subject history lie along what physicists call a timelike path, with the subjective successor of each event lying in the physical future of its predecessor. For subjects like us whose subject histories do happen to lie this way, the subject's viewpoint will move in a timelike direction on a path in spacetime as tenseless happening proceeds. For us, subjective time runs together with clock time, and to endure is to have some physical time pass. Yet despite this contingent connection between physical and subjective time, these two forms of time are logically distinct. Happening, which grounds subjective changes of tense, is a feature of spacetime as a whole and also is an inherent property of individual events. Physical time is a set of directions in spacetime, described by coordinates defined on spacetime with the help of clocks.

Minkowski versus Bergson: A Peace Proposal

Bergson is the philosopher best known for drawing the distinction between physical time and temporal flux.¹⁰ He portrayed time as primarily "duration," a qualitative, non-metrical principle which pervades experience but is ignored in scientific analyses of time. On Bergson's view, the physical and mathematical analysis of motion tends to lead to the false conclusion that time is simply a series of moments analogous to the points of space.¹¹ Quite different from Bergson's view is the standard interpretation of special relativity, primarily due to Minkowski.¹² On the latter view,

time is a dimension, or type of extension, rather closely analogous to space.¹³ People often feel that these two views of time are incompatible. Even if relativity theory does not strictly force one to believe that time consists simply of a series of time-positions or like items, it certainly makes such a view tempting.

The conflict between special relativity and the Bergsonian view of duration is real only if the Bergsonian position implies that time is solely a non-metrical principle, instead of a "spatialized," Minkowskian series of moments. (This would be the case if, for example, the physical universe were three-dimensional and contained objects which endure in Bergson's sense.) But the Bergsonian view of time as duration need not contradict the scientific view that time is a dimension of the world. The Bergsonian view still can hold if there can be duration *in addition to* a geometric dimension of physical time — that is, if the world is a four-dimensional spacetime in which real duration plays some role.

Bergson's own views, pushed to their logical conclusion, appear to imply something like this combined view. Bergson held that time is *cumulative* — that the past is contained in the present in a certain fashion.¹⁴ But if the past is contained in the present and the present really endures, then what really endures is the past and present together. If one takes seriously Bergson's ideas of duration and of the accumulation of the past in the present, then one is compelled to conclude that both the past and the present endure. Bergson also seems to have held that the present contains the future, though perhaps only as

"potentialities."¹⁵ If its claims about the cumulativeness of the past and the potentiality of the future are taken at face value, then Bergsonianism issues in the view that past, present, and future all endure as the present endures. In other words, Bergson's view of time implies that *it is spacetime or the spatiotemporal world which really endures*.¹⁶ This consequence agrees with my conclusion that the spatiotemporal world as a whole exemplifies happening. Such a view of spacetime reconciles the reality of spacetime with the reality of pure flux.

Chapter 10

Conscious Beings and Physical Things

Conscious subjects and physical objects are the most important constituents of the observable world. They are the most obvious, and probably the only, sorts of things which persist through time. Nevertheless, the existence of such things is not an obvious consequence of facts about how the world seems. The existence of momentary mental phenomena, and the apparent existence of physical objects, clearly follow from facts about how things seem. However, the existence of real subjects and objects, which really persist through time, may not follow as readily from these facts. (Recall my earlier discussions of Descartes and Hume.)

In this chapter, I will argue that certain facts about how things seem imply the existence of persisting conscious subjects. I will argue also that such facts imply the existence of real persisting objects of which subjects are aware, if we grant that some objects which seem to persist are real. Although I have discussed conscious subject histories in

some detail, I have so far avoided the question of the existence of conscious subjects as such. Here I will try to derive the existence of subjects with the help of subjective facts, and also will build the beginnings of an analogous case for the existence of physical objects.

Of course, one does not have to doubt the existence of subjects or objects to find these arguments interesting. Although a Cartesian or Humean skeptic could take these arguments as refutations of skepticism, my chief motive is not to refute skepticism but to find out whether certain features of the world (the existence of persisting subjects and objects) follow from facts about how things seem.

What Are Diachronic Objects?

We are aware of objects that seem to persist through time. Physical objects are the most obvious examples. A subject may experience what seems to be the same physical object for more than one moment in that subject's history. Thus, what appears to be a single physical object may exist for more than one successive consciousness event in a subject history. I will call an entity *subjectively diachronic* if it exists for more than one successive consciousness event in some subject history.

Physical objects often persist through time while being observed by subjects. Hence physical objects are subjectively diachronic items. Persons or selves also are subjectively diachronic items, at least if they undergo more

than one consciousness event.

Certain puzzles arise when we try to understand the commonsense notions of physical object and of conscious subject. Among these puzzles are the problem of personal identity, which I discussed in Chapter 5, and the similar problem of *physical identity*, which is the problem of how the stages in the history of a physical object form the history of a single, persisting object.¹ The history of a physical object, like that of a person, consists of a series of stages or of events (either one or the other depending upon the exact way one analyzes these histories). These stages or events are tied together in some way to form the history of a single object. It is somewhat problematical that some sets of stages or events are histories of objects while others are not. There also is a problem about our knowledge of temporally persisting objects: how can we be sure that there are any such objects at all? Philosophers have recognized that one does not need to doubt the reality of the physical world to doubt that there are persisting physical objects. The following discussion of some of these ideas may make this point clearer.

Some philosophers have held that the traditional concept of a physical object is not really necessary for understanding the world we discover through sense experience. On this view, the sensible world can be described completely in terms of entities other than physical objects, such as temporary stages or sense data.² If such a description is possible for any experience, then no set of experiences can serve as sufficient evidence for the existence of physical

objects. There always is an alternative description of the experienced world which does not make use of physical objects, and this description has a sort of simplicity on its side.³ If this view were right, then there still would be the strings of stages or events which we usually regard as the histories of physical objects, but there would be no physical objects of the kind in which people normally believe — that is, objects which really persist through time. We must ask whether physical events might frequently occur in temporal series which behave precisely as though they were the histories of diachronic objects, even if there is no single thing which exists during each of these histories.⁴

Philosophers also have expressed doubts about the existence of a persistent self underlying personal experience. Hume is well known for having such doubts.⁵ On views like Hume's, the stream of experiences that makes up the conscious life of a person suffices to account for all the facts of what we normally call personal experience; there is no need to postulate an enduring self behind the stream of experiences.⁶ Even the feeling that I am, that there is a self that persists through time, may be regarded as just another kind of experience in the stream.⁷ On this view there is no need of persons or subjects to account for any feature of sensory experience; histories alone will suffice. A more recent (and very different) critique of the traditional notion of a persisting self comes from Dennett, who has suggested that the self is a fictional entity of a specific sort.⁸

Doubts about the reality of temporally persisting objects

pose a real philosophical challenge. I am not speaking of doubts about the reality of the external world, but of doubts about the temporal persistence of objects in that world, such as I have described in the last two paragraphs. *Prima facie*, these doubts force one to choose between acquiescing in skepticism about persisting objects, or attempting to demonstrate the existence of persisting selves and physical objects.

One convenient way to overcome skepticism about a thing's existence is to find an object of whose existence you are fairly sure and then show that that object is the thing whose existence you doubt. To do this, you must show that the known object answers to the definition of the kind of thing you are skeptical about. The results of such identifications of the known with the unknown sometimes involve substantial side assumptions. (For example, Russell and Carnap attempted to portray physical objects as logical constructions.⁹ Such attempts will yield genuine objects, rather than convenient substitutes, only if one supposes that a physical object can be a logical construct of a certain sort — which I do not.) Fortunately, the theory of subjective time which I introduced earlier provides us with ready-made objects which fit the informal notion of *conscious subjects* in all important respects. I will spend the next several pages leading up to an identification of conscious subjects with those objects.

What Is a Persisting Object?

The history of a diachronic object, as we normally conceive of it, consists of a series of things that happen. Such a history normally includes, and may be composed entirely of, events. Mellor has discussed phases in the careers of objects, during which phases no change happens to the object. Such "thing-stages," Mellor claims, are best not regarded as events.¹⁰ Mellor may well be right about this, but his usage of the word "event" appears to be slightly different from mine.¹¹ Any thing-stage is an entity that transpires, that passes as time passes. It exemplifies happening, and hence is an event in my sense of that word. Thus we can say that the history of a temporally persisting object is a series of events.

We want to try to derive certain claims about the existence of diachronic objects from facts about how things seem. We should ask whether diachronic objects might be identical to anything whose existence we already have derived from those facts. Is there anything in the domain of consciousness events, as described in Chapters 2 through 9, that might be a diachronic object?

Some philosophers, most notably Broad, Russell and Quine, have proposed that diachronic objects are identical to, or should be identified with, their entire histories.¹² At first glance this identification looks implausible. The chief source of this implausibility is the fact that a diachronic

object exists as a whole at each moment in its history. Intuitively, if a diachronic object exists at a time, then at that time the entire object exists; all of its parts exist at that time.¹³ An object history lacks this feature; it has temporal parts,¹⁴ so there is no time at which the entire history presently is happening. Thus, the view that a diachronic object is a history implies that a thing with parts spread out over various times can exist *entirely at a single time*. This looks like a contradiction.

There are at least four sorts of things which a diachronic object might be:

(I) The most skeptical possibility is that it might not be anything. That is, there might not be any real diachronic objects underlying the flow of events. On this view, discourse about such objects is little more than a figure of speech. Hume took (or came close to) this position with regard to subjects.¹⁵ Bergson took a position at least reminiscent of it with regard to physical objects.¹⁶

(II) A diachronic object might be something distinct from its history — some sort of enduring substrate to which the object states or events in the object history somehow belong. This view is essentially the commonsense view of physical objects and persons: an object is not the same thing as a history, but it *has* a history. Common sense suggests that where there is a physical object history there also is something else (a physical object), and that this "something

else" exists entirely at each moment of its history and persists from one moment to the next. One can have a theory of physical identity without knowing exactly what this "something else" really is.

This view seems natural in the light of certain philosophical views of physical and personal identity which do not depend upon the identification of diachronic objects with histories or with composites of events or stages.¹⁷ These views of identity can be carried through even if we do not identify an object with its history; they demand only that the existence of a suitable history implies the existence of an object of a certain sort. Such theories provide analyses of physical or personal identity, but need not explicitly say what kind of an entity a physical object or a person really is.

(III) A diachronic object might be something which is real, but which only seems to exist entirely at each moment. This way misses the notion of a diachronic object. To be diachronic is to persist from one moment to the next. If an object is diachronic, then somehow or other it persists as a whole from one moment to the next. If an object does not exist in this way then the object is not diachronic. This alternative also misses the informal notions of physical object and of person. Things which never exist in their entireties are not normally regarded as persons or as physical objects.¹⁸

(IV) A diachronic object might be an entity which exists entirely at each moment of its existence, *yet also happens*.

That is, it might be something which has temporal parts located at various times, and yet exists entirely at each moment in its history. This alternative seems self-contradictory, but it is inconsistent only if one assumes that temporal parts which exist at one time cannot also exist at another time. This assumption seems obvious on the commonsense conception of time, and on most philosophical accounts of time as well. However, we have arrived at some new views of time in the preceding chapters. We should consider the possibility that these views may make the assumption just mentioned less obvious.

We need to look closely at alternative (IV). On its face, (IV) is simply inconsistent; it appears to imply that there can be an object which exists entirely in the present, and yet has parts which exist in the future or in the past. But on the view of time which I presented earlier, this is not a contradiction at all. The following example shows why not.

Consider a coffee cup which exists at 1 pm. Suppose that the cup has a part — call it *x* — which is an event. The part *x* is located at a time other than 1 pm — say, at 1:30 pm, during a later stage of the cup's existence. We have seen (Chapter 9) that all events both exist and happen tenselessly. Thus it is the case at 1 pm that *x* both exists and happens, even though *x* is not present at 1 pm. Of course, *x* is not located at the time 1 pm, but since *x* exists and happens tenselessly, it is correct to say, at 1 pm, that *x* both exists and happens. (If this sounds strange, re-read the discussion of tenseless happening in Chapter 9.)

This argument extends to any object which has events as parts. Therefore, if an object has events as parts, that object may have parts which are not all located at the same moment of time, and yet still may *exist tenselessly* as a whole at each moment in its existence. This is possible because every event exists tenselessly.

Thus, alternative (IV) is not as inconsistent as it looks. An object which has parts located at different times still can exist as a whole at each moment of its history. At each moment, all of the object's parts really exist. They are located at different times, but nevertheless each of them exists at each moment, in the tenseless way described in Chapter 9. The possibility of all this depends upon the distinction between an event's existing at a time and its being present at that time. If these two notions are not distinguished (and in ordinary discourse they seldom are), then the claim that an item present at one time may be part of an object at another time leads to contradictions.

Of the four alternatives listed above, only (IV) and (II) promise to give us an object which really might be a subject or a physical object. (I) amounts to abandonment of belief in diachronic objects, while (III) implies that there are no genuine diachronic objects, though there may be pseudo diachronic objects. (II) yields an object, but there is nothing we have discovered in our project that could be such an object; for all we know, (II) might be true, but if it were, we could not determine this from the conclusions we have drawn so far. Only (IV) allows us to preserve any immediate hope of deriving the existence of a diachronic

object within our present conceptual scheme. If (IV) holds, we might be able to identify diachronic objects with their *histories*. We already know that there are subject histories; if diachronic objects are histories, this would allow us to conclude that there are subjects. Perhaps we can do something similar for physical objects.

The above arguments suggest that some histories are diachronic objects. However, they do not show that diachronic objects are histories. The conclusion that histories of certain sorts are diachronic objects does not imply that the subjects or physical objects which we think we experience are just histories. Perhaps alternative (II) is right, and subjects and physical objects exist but are not histories. Perhaps physical objects and conscious subjects are real, but the arguments presented in the previous chapters simply leave them out, and do not allow us to show that they exist. These entities might have to be secured in some other way, or (alternative (I)) they might not exist at all.

These other possibilities can be disposed of or rendered harmless. I will do this next.

Intuitively, one thinks of a physical object as a unified whole which has a history and which exists at every moment in that history. One witnesses a stream of events of a certain sort (such as light reflection happening here, then there, then somewhere else), and one supposes that each of these events is part of the history of a single underlying thing. According to our intuitive understanding of physical objects, the physical object is the thing to which the events of the object

history happen — the substratum in which all those events inhere. But the preceding arguments show that the physical object history *itself* is a substratum of precisely this sort. This history exists at each moment of the physical object history, and all events in that history take place within it. And each event in the physical object history is, from one point of view, *something that happens to the history*. To recognize this last point, imagine that the object is a billiard ball and its momentum is changed by the impact of a cue at time t . Then the trajectory of the ball is changed at time t ; hence at time t , something is added to the ball's history. If the ball is destroyed at some time t' , then the history is supplied with an end at t' . Each event happening to the ball takes place in the history and is something that happens to the history; the fact that the event also is a part of the history does not change this. Hence a physical object history has the essential characteristics which define the object *undergoing* a history. A physical object history has properties which we expect to find in the physical object whose history it is.

Similar remarks apply to subject histories. Intuitively, a subject is the locus of consciousness; it is the entity in which consciousness resides. It is an entity for which facts can seem to be the case. A subject history made of consciousness events answers nicely to this description. It is an entity for which facts can seem to be the case (with different facts being the case for different consciousness events in the history). It is a locus of consciousness. It is a thing which has consciousness in it; consciousness is seated

in it. It has a subjective world of apparent facts and beings at each moment. In a certain obvious sense, it is "made of consciousness."

If the conclusions about time discussed in Chapters 7-9 are right, then physical object histories and subject histories answer to the informal notions of physical object and of subject respectively, in all important respects. Hence there is no excuse not to identify these objects and subjects with their histories. Once we have done this, there is no need to doubt the existence of persisting physical objects or of subjects as long as streams of events of the appropriate kinds exist.

We do not need to postulate a separate underlying persisting thing or substratum, distinct from its own history, just to secure the existence of a diachronic object. *The histories already are persisting things, and already are substrata for events.* We cannot strictly rule out the possibility that there is a persisting object distinct from the history. But if there is no such object then there still is a diachronic object — namely, the history itself.

This conclusion completely defeats any skepticism about the existence of diachronic objects. If a follower of Hume's skepticism challenges us to show that a self exists amidst the stream of "impressions,"¹⁹ we can reply that even if there is no self over and above the stream of impressions, the stream itself still is a diachronic object and has the most essential features of what normally is called the self. That is, the stream of impressions to which Hume points *is* the very self for which Hume asks. Skepticism about the existence of

persisting physical objects meets with a like fate. If there is no physical object distinct from the stream of physical events or stages which make up a physical object history, then that history *is* the physical object.

One objection to alternative (IV) is that it seems to imply that a diachronic object exists at times earlier than its beginning or later than its ending in time. If an object exists tenselessly, then at each time it is correct to say that that object exists. Hence a diachronic object exists as a whole at each time, even before its history begins or after its history ends. This seems to imply the falsehood that the object has no beginning or ending in time. This objection, like other objections to alternative (IV), is based on a misunderstanding of the distinction between existing at a time and being located at a time. The tenseless existence of a diachronic object does not imply that that object always is present. The object is temporally located between the beginning of its history and the end of its history; it can be present only during those times. Nevertheless, it exists tenselessly. At times before or after its history, it is located in the future or in the past instead of in the present. It exists — but that does not make it present.

All objections aside, the view that physical objects and subjects are histories may simply seem strange. Normally we do not think of ourselves or of physical objects as things that *happen* in the same way that dawns and fires happen. The following argument may help to counteract this uneasiness. It suggests that the persisting objects which we perceive actually *appear to us to be temporally extended*,

whether or not we realize it or think about it.

Suppose that a coffee cup is subjectively diachronic for someone's subject history. The cup exists for each consciousness event in a certain segment of that history. Think about two consecutive consciousness events in that segment. Call those events *y* and *z*. The cup exists for *z*; hence for *z*, the cup is present. But since *y* is continued in *z*, *z* may also involve the appearance that the cup was present a moment before. (Recall what I said in Chapter 4 about hearing a loud noise.) Thus, the subject can notice, during *z*, that the cup exists in the immediate past. But at the moment *z*, the cup also is located in the present moment of subjective time. These two facts together imply that the coffee cup is temporally extended in subjective time — that is, it occupies at least two subjective moments, as viewed from the vantage point of the single subjective moment *z*. Thus, the experience of the coffee cup involves items at different subjective moments being present simultaneously to a subject's consciousness. In other words, the cup *as experienced now* has temporal parts at different subjective times.

This argument, when suitably generalized, suggests that all subjectively diachronic objects in our experience have temporal parts which exist at different subjective moments. This supports the view that diachronic objects are histories of some sort. The argument is suggestive rather than conclusive, but at least it shows that it is not too strange to suppose that the diachronic objects which we perceive are temporally extended. Our perceptions of physical objects

carry over into future moments in a way which allows us to be aware of those objects (directly or indirectly) at two different subjective times at once.

In any case, we achieve no conceptual economy or simplicity by refusing to identify diachronic objects with their histories. If we postulate, along with each diachronic object history, another object which we identify as the diachronic object itself, then we are positing two entities of different kinds rather than one entity. In view of what I just said about the coffee cup, this other object probably would have to be temporally extended anyhow. Such redoubling of objects is not inconsistent; we might find other reasons for duplicating entities in this way, but economy is not one of them.

Why Things Appear to Be Temporally Unextended

A physical object is a history. Nevertheless, physical objects appear to us to be purely spatial and devoid of temporal parts. The physical basis of this fact is well-known to physicists: at any given time in an observer's history, the observer can perceive only a three-dimensional slice of a four-dimensional object history. In other words, at any given moment in one's own history, one only perceives a brief stage in the object's history. There is little more to be said about the reason why things appear three-dimensional. In this section I will offer a few remarks connecting this

reason to the ideas presented earlier in this chapter.

Look at a coffee cup. At any moment (consciousness event) in your history during the time you are looking at the cup, only a short stretch of the cup's history exists for you. This stretch amounts to a three-dimensional slice through the history of the cup, which is four-dimensional. (This slice has finite temporal thickness, since our sense organs obviously do not have infinite time resolution.) But according to the arguments of this chapter, the history of the cup *is* the cup. Hence if spacetime is four-dimensional, then a coffee cup, as it really is, is four-dimensional. It only seems to be three-dimensional. At any moment of your subjective time, you perceive the cup as an instantaneous (three-dimensional) object. These remarks apply not only to coffee cups, but to all physical objects which you might perceive. At any given time, your senses (actually, stages in the histories of your sense organs) receive information from the events in a brief slice of the object's history and not from other events in the history.

A diachronic object is "wholly spatial" in a generalized sense: at any given time, it is the case that all parts of the object exist. Hence there is an extended but legitimate sense in which a physical object is wholly spatial or three-dimensional. But speaking literally, we can say that diachronic objects have more than three dimensions, because they are temporally extended histories.

It is important to note that I have not yet showed that the existence of persisting physical objects follows from facts about how things seem. Instead, I have shown that

persisting physical objects are histories if they really exist. The argument here also shows that the appearance of a physical object is an apparent history — that is, to have an experience of a physical object (whether illusory or not) is to have an experience of a history.

Closing Remarks on the Self

The view of the persisting conscious subject presented in this chapter allows us to overcome any lingering skepticism about the self. There can be a self even if there is no diachronic ego of the traditional sort — no core of a person's makeup which exists first wholly at one moment, then wholly at another moment, and so forth.

My account of subject history leaves open the possibility that there is nothing to a subject but a stream of consciousness events. Certainly such a stream of consciousness suffices to account for the properties of personal experience. But according to my view of time, such a stream of consciousness also is a diachronic entity. One consequence of this identification of subject with history is that the self not only exists now, but also contains its past and future (or possible futures) as parts. This conception of the subject resembles Bergson's in some respects. Bergson held that for living beings, the past persists and builds up in the present; apparently he also held that the future exists as "potentialities."²⁰

Another consequence of this account of the self is that no

subject is located exclusively at a particular time, even if that subject does pass through stages at particular times. Hence subjects transcend time in a way reminiscent of the way in which the personal idealists George Howison and Borden Bowne said they did.²¹ Despite their possession of temporal lives, subjects are in one sense timeless. This transtemporal character of the subject is fully compatible with the subject's possession of a life in time.

Of course, the consciousness of a subject can be confined to the present moment in an obvious sense, even though the subject is temporally extended. During a particular consciousness event in a subject's history, that consciousness event is the only consciousness which the subject presently has. At that subjective moment, the current consciousness event is the subject's only present consciousness. From the subject's point of view, this consciousness moves forward in time as one consciousness event seems to replace another and to become the present consciousness of the subject.

From Brain to Cosmos

Chapter 11

The Structure of the Self

In the previous chapter I proposed the beginnings of a theory of the conscious subject. In this chapter I will investigate the *psychology* of conscious subjects — that is, the specific mental characteristics which subjects must, or may, have. I will concentrate on two characteristics which are important in human mental life: the presence of unconscious mental processes and the possibility of a lack of unity within the self. Both of these features of mental life play important roles in psychological and philosophical theories about the human mind.

The importance of the idea of the unconscious in psychological theories is well-known. The hypothesis of a nonunified self or mind also has psychological and philosophical importance; it plays major roles in theories of irrationality¹ and of certain problems of personal identity.² My chief aim here is to show that the existence of unconscious contents and of multiple "compartments"³ within the ego is compatible with my account of the subject.

Indeed, my account predicts that these features are possible for subjects. Another aim of this chapter is to show that disunity of the ego, of the sorts postulated by the theories mentioned above, cannot compromise in any way the unity of the conscious subject. It is wrong to use the psychological disunity of the subject to infer that there is not really one single subject.

The Conscious Subject: A Review

Let me begin by reviewing in three paragraphs some of my earlier conclusions about conscious subjects.

A subject's history consists of a series of consciousness events. These events are specious moments of conscious life; they are not to be confused with individual events of sensing, thinking, feeling, and the like, since one can sense, think, feel, etc. simultaneously during the course of a single consciousness event. Also, the same subjective mental process — the same thought, emotion, or sensation — may go on in you during more than one consciousness event. (For example, you may look at a green rectangle and continue to stare at it for a while, thereby having a single prolonged sensation of green.)

Subjective time is a feature of the stream of consciousness events; this stream constitutes the history of a subject. The subjectively experienced succession of events results from the subjective existence, for each consciousness event, of a consciousness event which preceded it. There

are subjective tenses and there is a subjective movement of time; these can exist for the consciousness of the subject, although tense is not objectively real.

A subject contains as parts its own past and future. The subject's past and future exist objectively, and the subject's own immediate past exists for the subject's present consciousness event as well. This inclusion of the past and the future in the subject make the subject a temporally extended entity. Due to certain features of subjective time, this temporally extended entity also is an entity which *persists* through time. There may not be any persistent, substantial self apart from this entity.

A Subject Is Not an Ego

Some philosophical views of personality seem to imply that a person is an ego, or conscious mind. Descartes is notorious for making this identification explicit — or seeming to do so. The soul with which Descartes identifies a person is "a thinking being"⁴ and is more or less transparent to introspection⁵; hence it is, more or less, what we today would call an ego. (The well-known fact that Descartes conceived this soul as a nonphysical substance distinct from the brain is irrelevant here. The relevant fact is that on Descartes' view, a conscious subject is first and foremost a being with intellect and volition. One need not accept Descartes' dualism to accept this.) Views like Descartes' suggest that drives and similar psychological

processes do not really belong to the self. On such views, these processes might be regarded as bodily rather than as mental — perhaps as disturbing influences on the self arising from processes in the body.⁶

Other philosophers, before and after Descartes, have rejected this ego-oriented viewpoint. They have developed views of the self in which the rational, moral ego does not occupy the entire picture. Plato admitted "appetitive" and "spirited" factors (that is, urges, drives and emotions) into the makeup of the soul.⁷ Some German idealists, notably Hartmann and Schopenhauer, stressed the ascendancy of unconscious over conscious processes.⁸ More recently, several philosophers have emphasized the possible *disunity* of the ego. Donald Davidson and David Pears both have explored models in which different mental components or compartments exist together within the same person.⁹ Some scientists and philosophers have interpreted neurosurgical findings in ways that seem to imply that the ego can become disunited under certain circumstances.¹⁰ For quite different reasons, some current philosophers of mind portray the human mind as arising from interacting "homunculi."¹¹

The view that one's entire self consists solely of one's ego is hard to reconcile with most modern psychological theories. Psychologists long have known that there are "mental" processes going on in us of which we know nothing. Psychoanalytic theory depicts the unconscious as the largest part of the psyche.¹² But even those who reject

psychoanalysis have to recognize that many psychological phenomena happen without the participation of what we call our consciousness. Even a person who knows nothing of psychological theories cannot help but notice that people undergo dreams, involuntary habits, and other phenomena which seem to originate from the person but not from the person's conscious self.

The psychological theories and observations mentioned in the last paragraph do not strictly force us to believe that unconscious or non-rational psychological processes take place within the self, rather than originating from some (bodily) source outside the self. One can acknowledge the existence of unconscious phenomena, and even of the psychoanalytical unconscious, and still tenaciously hold that one is one's ego and nothing more. One can stubbornly hold that reasoning is an activity of the self while also maintaining that the drives are not activities of the self. Nevertheless, the psychological facts do suggest that the equation of the self with the ego does not quite balance. Although that equation does not really contradict any of the data of psychology, it is difficult to find any supportable basis for such an identification. From a metaphysical point of view, "bodily" drives are just as "mental" as abstract thoughts or religious feelings. All these processes take place in the realm of subjective experience, even though some of them feel more "physical" than others. The drives either are conscious mental contents or are capable of "surfacing" and becoming conscious mental contents. Even if one attributes the causation of the drives to processes outside the self, one

cannot help but regard them as parts of the life of the self once those causes have acted. The same conclusion applies to alleged unconscious things that are even more difficult to make conscious, such as repressed fears and information denied through self-deception.

A subject's history is a stream of consciousness events. Each of the consciousness events may involve many different subjective processes; for example, a thought, an emotion, a biological urge, and a visual sensation may all play roles in the same moment of subjective life. Each of these subjective processes contributes to the consciousness event as much as any of the other processes does. Each process plays a part in experience in exactly the same way as do all the others. All of the processes play parts in the subject's "inner world" — the mental realm comprised by the way things seem. There are no grounds for regarding some of the contents as "truly mental" and others as exclusively "bodily," or for regarding some as belonging to consciousness and others as somehow extraneous to the mind.

Unknown Experiences

The following argument points out a way in which a subject may undergo an experience *without being able to know that that experience has happened*. The conclusion that this is possible has an interesting consequence: that some of the *conscious* processes going on within a subject

may not be directly knowable by that subject.¹³

Suppose that a subject John has a consciousness event — call it *x* — and that one of the psychological contents which exists for *x* is (to use a Freudian-sounding example) a wish to commit adultery. If *x* is not the last consciousness event in John's history, then there is another consciousness event — call it *y* — in John's history such that *x* exists for *y*. But this does not imply that it is the case for *y* that a wish to commit adultery has occurred. This lack of implication follows from the logical incompleteness of consciousness events, which I discussed in Chapter 3. It is the case for *y* that *x* exists, and it is the case objectively that an adulterous wish is among the contents of *x*. But this does not imply that for *y*, it is the case that an adulterous wish occurs. Thus it is logically possible that it is *not* the case that for *y*, an adulterous wish occurs. The logical incompleteness of consciousness events ensures that a consciousness event may "overlook" some of the contents of its predecessor in this way. Such "overlooking" leaves the subject *incapable of thinking about some of his or her own contents*.

In Chapter 4 I argued that the contents of a consciousness event which exists for another consciousness event can be known with certainty under certain conditions. The argument of the preceding paragraph does not contradict this. If *y* involves continuance of *x* and *y* also involves feeling the adulterous wish present to *x*, then John might be able to know with certainty that the wish had occurred. But the fact that *y* involves continuance of *x* does not imply that *y* involves awareness of *all* of the content of *x*. It could be

the case that the existence of *x* is a fact for *y*, but the presence of the wish is not. Continuance of a consciousness event *x* need not involve awareness of all subjective facts about *x*. There is no reason why it must involve awareness that a particular fact is the case for *x*, even if that fact really is the case for *x*.

The Nature of the Unconscious: An Hypothesis

This possibility of having an experience and not being aware later that one has had that experience amounts to the possibility of unconscious mental processes of a sort. In the adultery example above, John never notices directly the presence of his own adulterous wish. It is conceivable that John never becomes aware of the existence of that wish at any later subjective time. If that were the case, then John never would be able to think consciously about the fact that he once had the adulterous wish. Nevertheless, he really did experience that wish when he had it. John's adulterous wish is much like what psychologists call an item of unconscious content. It is genuinely mental, an element of John's inner life; if we like, we can say that it is in John's psyche, just as much as John's conscious thoughts are in his psyche. Yet the wish, though technically a content of consciousness, is *unconscious* in the sense that John, who has it, never knows that it happened.

Contents which become unknowable to the subject in this way are genuine contents of a subject's consciousness. They

are in the subjective realms of the consciousness events of a subject; they are things that exist, or are the case, for that subject's consciousness events. But those contents are unconscious in the sense that the subject who has them does not know that they are there. In particular, they never can be discovered through introspection. One must be careful not to fall into verbal confusion here, for a content that a psychologist would call "unconscious" still can be "conscious" in my sense of that word. (Recall my discussion of consciousness in Chapter 2.) Unconscious contents of this sort are facts and things that one experiences without knowing that one has experienced them.¹⁴ They belong to the subject in exactly the same way in which, for example, a well-considered thought about morality belongs to the subject.

The chief difference between the unconscious adulterous urge and the conscious moral thought lies in the way in which these two elements of experience are experienced *after* they occur. The presence of the unconscious content involves the subjective truth of some fact for the subject. Yet it never is the case for the subject that that fact *once was* the case for that subject. The consciousness event for which that fact is the case nevertheless exists for the consciousness of the subject.

The conception of the conscious subject which I put forth in earlier chapters has turned out to be far richer than it first appeared to be. The stream of consciousness events that constitutes a subject's life need not contain only psychological processes which we ordinarily regard as

"conscious." It also may include processes which a psychologist would call "unconscious."

The preceding discussion of unconscious processes shows that I have been using the term "consciousness" in a sense different from the standard psychological one. This should be no surprise; in Chapter 2 I said that I would use the word "consciousness" in a special sense. There I characterized consciousness as the possession of a way that things seem, or of a point of view. Phenomena that are unconscious in the psychological sense can still be contents of consciousness in my sense. Psychologists and others often use the word "conscious" to refer to the realm of psychological phenomena of which the subject *can know directly*, that is, to the contents of the subject's ego; they designate all other mental processes "unconscious."¹⁵ This is not the sense of "conscious" which I laid out in Chapter 2. My characterization of consciousness was meant to capture certain commonsensical and philosophical usages of "consciousness" — not standard psychological usages. It encompasses all subjective phenomena, whether known to the subject or not. It is a concept far broader than what many psychologists and others call "consciousness."

Conscious mental life, according to the sense of "conscious" in which I am interested here, may encompass much or all of what some psychologists call the "psyche" — the entire inner or psychological life of a subject.¹⁶ We may speculate that the unconscious mind is just that section of the subject's consciousness which is unknown to the subject for reasons described above. If this speculation is correct,

then one's unconscious might be a genuine part of one's self, and indeed the *greater* part of that self.

Of course, this is only a speculation. We have not shown that real unconscious mental phenomena are phenomena of the sort which I have described. However, this interpretation of the unconscious appears to be compatible with all that we know about the unconscious.

It is important to keep in mind that this hypothesis is not a full explanation of the unconscious, or of any particular unconscious mental phenomenon. An explanation of unconscious mental life would have to account for the phenomena of that life in terms of brain activity, as well as dealing adequately with the metaphysics of those phenomena. My remarks about the unconscious are meant only to make one point about the metaphysics of the unconscious. This point is that unconscious processes may be, in a wider sense, *conscious* — that is, they may be conscious processes of a peculiar sort. They may involve subjective fact and a viewpoint, just as do the phenomena which are more conspicuously conscious.

A conscious subject has a subjective side as well as an objective side — an inner as well as an outer aspect. There are ways things seem to a subject, and there are apparent realms of things that exist for a subject (or at least of facts that are the case for the subject) in addition to the world of things that exist objectively. It is this involvement in subjective fact that makes a conscious subject a truly psychological being — a being which has a psychological life. It is subjective fact which confers upon psychological

phenomena their peculiarly mental character. The occurrence of subjective fact in the world is not restricted to the sort of consciousness demarcated by the boundaries of our egos. Instead, it pervades all genuine psychological life. All beings endowed with real subjective life of any sort, conscious or unconscious, are beings for which facts sometimes are the case.

The Absolute Unity of the Subject

The idea of a non-unified mind is contrary to some views of human nature. However, as I have mentioned, this idea is important in the philosophy of mind and of action as well as in psychoanalytic theory. As I will show, my account of subjects allows subjects to be psychologically fragmented or compartmentalized in a way which seems adequate for the purposes of philosophy and psychology. However, this fragmentation does not compromise the underlying unity of the subject in any way. The subject still is one stream of consciousness. It is merely the *content of the subject's experience* which is compartmented. There are no divided selves; there are only selves with divided experience.

We may characterize a psychologically fragmented subject informally as one who has several distinct contents of which he/she is aware, but who is not aware of all of those contents together. For example, a psychologically fragmented person might love and hate the same person, might feel the love and simultaneously feel the hate, and still

fail to notice that he/she both loves and hates the same person at the same time. On some philosophical accounts of irrationality¹⁷, mental contents are grouped into compartments in such a way that contents in different compartments do not affect each other's behavioral roles in the usual way and to the usual extent. For example, suppose that a person knows enough to realize that a belief P is irrational but believes P anyhow; on a view explored by Pears, this may indicate that P and the "cautionary belief" that P is irrational are in different "systems."¹⁸

The account of the conscious subject presented in this book suggests a way to understand this compartmentation. Let us begin with an example. Suppose that a subject (called A) has a consciousness event x whose contents include love of another person (called B) and hate toward the same person. Let y be the successor of x in A's history. Suppose that y involves direct awareness of both the love of B and the hate toward B — that is, for y love of B is present, and for y hate toward B is present. Then A is aware of love of B and of hate toward B. Yet because consciousness events are logically incomplete, it is possible that A never has any direct awareness of the love of B and of the hate toward B *together*.

The Logic of Confusion: Some Technical Points

Let us look at this example in great detail, using the language of subjective fact.

If at y , A both loves and hates B, then the following are true:

- (1) For y , love of B is present.
- (2) For y , hate toward B is present.

If A becomes aware that he simultaneously loves and hates B, then the following is the case:

- (3) For some consciousness event z in A's life, love of B is present and hate toward B is present.

Note that (3) does not follow from (1) and (2). Due to the possible logical incompleteness of the consciousness event y in (1) and (2), we cannot pass from (1) and (2) to

- (4) For y , love of B is present and hate toward B is present.

We also cannot infer (4) with y replaced by some other consciousness event besides y , since (1) and (2) only tell us what is the case for y .

We cannot pass from (1) and (2) to (3) because consciousness events are logically incomplete. A logical consequence of what is the case for a consciousness event need not be the case for that consciousness event. In this instance, both "Love of B is present" and "Hate toward B is present" are the case for y ; these two sentences imply the conjunction "Love of B is present and hate toward B is

present." But this does not allow us to infer that the conjunction is the case for y . That is, it need not be the case that A is aware that he both loves and hates B . Because y is logically incomplete, y may involve awareness of love of B and awareness of hate toward B , but still need not involve awareness of the combined love-and-hate toward B . The two emotional tones may be noticed separately without ever being caught together.

It may seem odd that a basic logical principle such as the rule of adjunction (from L and H , to infer $L\&H$) may fail to hold for our awareness in this manner. Actually, such failures are commonplace — and needless to say, they do not involve any real contradiction. The following example illustrates one such failure. I glance briefly at a disorderly pattern of eleven dots of different colors. At once I am aware that the dots are distinct from each other; they are at different locations in my visual field, and they are differently colored. But the fact that there are 11 dots may not be a fact of which I am immediately aware; I may have to count to find out that I have in fact seen 11 dots. In this example, all the dots initially are perceived in one perceptual act; it may well be that all the dots are perceived during a single consciousness event E . Suppose, for the sake of the example, that this is what happens. If we give the names ' x_1 ', ' x_2 ', ..., ' x_{11} ' to the dots, then the following premises are the case for E :

(P_1) x_1 is a dot and is not identical to any of x_2, x_3, \dots ,

From Brain to Cosmos

x_{11} .

(P₂) x_2 is a dot and is not identical to any of $x_1, x_3, x_4,$
..., x_{11} .

(P₃) x_3 is a dot and is not identical to any of $x_1, x_2, x_4,$
..., x_{11} .

...

(P₁₁) x_{11} is a dot and is not identical to any of $x_1, x_2, x_3,$
..., x_{10} .

The conclusion that there are precisely 11 dots follows from these 11 premises by logic alone. (There even is a way to express "there are 11 dots" in the language of quantifier logic alone, without using any numbers.¹⁹) All of the 11 premises of this argument are true for my consciousness at the moment that I see the dots, provided that all the dots are in my visual field and that my visual capacities are normal. (Note that this does *not* mean that I have named the dots, whether with names like ' x_1 ' or otherwise.) Yet the conclusion which follows from these premises — that there are 11 dots — is not the case for E. To confirm this conclusion, I have to count. Although it is not the case for me then that there are 11 dots, neither is it the case for me that there are *not* 11 dots. Nevertheless, it is the case objectively that either there are 11 dots or there are not 11

dots.

The fragmentation of our consciousness of our emotions provides an example of logical incompleteness analogous to the example of the eleven dots. This incompleteness is reflected in the psychological fact that feeling love and simultaneously feeling hate does not guarantee an awareness that one is feeling love and hate at the same time. It is possible to experience love and simultaneously to experience hate without experiencing the combination of the two — that is, without noticing that one has conflicting feelings.

There is another way in which two feelings may be kept from simultaneous awareness. This way is closely analogous to the way in which contents may become unconscious. A person can fail to be aware, at any given time, of one of the two opposing feelings — but the person may be unaware of one feeling at one time and of its opposite at another time. A consciousness event in A's history may have hate toward B but not love of B as contents, while another consciousness event may have only the love, not the hate. Thus A is aware of loving B or of hating B, but never is aware of both at once. In this form of fragmentation, one content is "conscious" at a given moment; the other content is "unconscious" at that moment. At a particular time, A may only love B or only hate B. If events of the two kinds alternate rapidly, A may appear to other subjects to love and hate B at almost the same moment. A may or may not become aware of these alterations.

The logical incompleteness of consciousness events allows mix-ups like these to happen with beliefs as well as with emotions. A subject may fail to see that there is a contradiction among contradictory beliefs, even though the contradiction is glaring. Psychologists describe certain real-life cases of this as the possession of "logic-tight compartments."²⁰ Such compartmentalization can be understood in much the same ways that the having of opposing feelings can be understood. A subject A who holds contradictory beliefs that P and that Q might have a consciousness event x during which the belief that P and the belief that Q are present. Yet he might have no consciousness event in which he is aware that he believes both P and Q. If this happens, then A is unaware that he has both beliefs. He cannot convict himself of inconsistency, in spite of his inconsistency.

The second sort of fragmentation described above — in which two feelings become conscious but not at the same time — also may happen with beliefs. One need only replace the love of B and hate toward B in my previous example with the belief that P and the belief that Q.

These simple examples extend easily to more complicated cases in which a compartment "contains" several beliefs, feelings, or the like. One can get complex compartmentalizations in this way. For the record, I will write out in full one example of this sort.

Suppose that the following statements are true of a consciousness event y in a subject's history. As before, let P and Q be two mutually contradictory sentences (or

propositions) and let B be a person.

For y, hate toward B is present.

For y, acceptance of P is present.

For y, hate toward B is present and acceptance of P is present.

For y, love of B is present.

For y, acceptance of Q is present.

For y, love of B is present and acceptance of Q is present.

It is *not* the case that for y, hate toward B is present and love of B is present.

It is *not* the case that for y, acceptance of P is present and acceptance of Q is present.

It is *not* the case that for y, hate toward B is present and acceptance of Q is present.

It is *not* the case that for y, love of B is present and acceptance of P is present.

In other words: for y, "Hate toward B is present" and

"Acceptance of P is present" are the case together, and so are "Love of B is present" and "Acceptance of Q is present"; but the hate toward B and the acceptance of P cannot coexist with the love of B or the acceptance of Q. The existence of such a consciousness event would imply a mental situation in which the hate toward B and the belief that P are in one "compartment" and simultaneously the love of B and the belief that Q are in another "compartment." Worse yet, one can imagine an example in which these two compartments enter the subject's awareness at different times.

It should be clear by now that my model of the subject allows for the possibilities of fragmentation and compartmentation of the self. These possibilities are consequences of the logical incompleteness of consciousness events. (Note that I have not provided any kind of an *explanation*, neurophysiological or otherwise, for the fragmentation or compartmentation of the self. Rather, I have simply shown that my account of the subject allows this possibility.)

The Self: Real Unity in Apparent Multiplicity

Now we can see how a human subject might suffer from psychological disunity and fragmentation even while that subject remains *absolutely one* in an ontological sense (that is, remains a single subject with a single stream of consciousness). A subject's inner life may contain incompatible feelings, thoughts, and impulses. It may have

many compartments. It may be subject to rapid changes and fluctuations of feeling and of belief. A subject might even be quite contradictory and manifold in all of these ways without ever being aware of his, her or its own fragmentation. Yet such fragmentation of experience does not change the fact that the subject is *one entity* — a *metaphysical unity*. At each moment of a subject's life, there is one underlying stream of awareness, one diachronic seat or substratum of awareness.

Psychological disunity does not imply the metaphysical disunity of the self. Psychological unity actually has little to do with metaphysical unity. It is not the self, but the content of the experience had by the self, which can fail to be one.

The analysis of psychological disunity presented in this chapter shows that philosophical viewpoints which uphold the unity of the self are compatible with the hypotheses that the mind is partly unconscious and suffers from fragmentation and impermanence. The idea that the self is a true individual cannot be threatened by the disunity of our experience.²¹ Of course, the disunity of the self can have practical consequences; it may explain weakness of will and so forth. But even a weak-willed subject is an ontological unit. In such a subject, the unity may be well-hidden, but it is there.

A subject may have different conflicting thoughts, feelings, and intentions at the same moment. A subject also may have different and conflicting contents at different moments. Behind this panorama of psychological fragmentation is an underlying metaphysical unity — the

From Brain to Cosmos

diachronic subject, tenselessly real, *possessor* of the inner realms in which the fragmentation occurs. There is at all times a single flow of consciousness which belongs to the subject. This stream, which exists both temporally and tenselessly, is the real conscious subject, who may be far more inclusive than the empirical self one normally thinks one is. The subject is a true *individual*.

Chapter 12

Personal Identity: Some Problems

In this chapter I will explore a specific application of the theory of conscious subjects developed earlier. I will address the puzzling question of what happens when one conscious subject divides in two, or when two conscious subjects merge into one.¹ Problems about dividing and merging subjects have been studied extensively by philosophers. Such problems may be of great practical significance. In neurosurgery, cases arise in which a person's brain is divided down the middle, resulting in a person who seems in some respects to be of two minds.² Attempts to understand the philosophical issues raised by such cases can lead to difficult logical puzzles; the toughest puzzle cases regarding personal identity are those in which persons divide or fuse.³

The problems of splitting and merging persons are among the standard philosophical problems about personal identity. Here I will study the same problems, but with conscious subjects, as defined in Chapters 5 and 10, taking the place of

persons. (The difference between the notion of person and that of conscious subject was discussed in Chapter 5.) The ideas about subjects and time developed in earlier chapters will enable me to propose solutions to these problems. These solutions depend crucially upon my earlier conclusion that a subject is, or is intimately related to, a history of a special sort. I will show that if one sticks to my view of the subject at every step in the analysis of dividing and merging subjects, the problems about such subjects begin to look very different and to be much more tractable.

Splitting Subjects

Consider first the case in which one subject splits to become two subjects. This is the most realistic of all the puzzle cases I will discuss, since it parallels what happens during split-brain surgery for epilepsy.⁴ (Later I will argue that this parallel is not close enough to justify regarding a split-brain patient as a double subject. But the parallel might be exact if the patient were split into two nearly equal, *fully separated* parts. Such a split eventually might become technically possible.)

To understand what happens during the splitting of a conscious subject, we must study in detail the different things and events that appear or disappear during this remarkable process. It is much easier to do this if we label those things and events with symbols. I will do this here, though I will not be using any symbolic logic or other

mathematics. The following argument is somewhat intricate, but intricacy is unavoidable when one is dealing with an intricate problem.

Let *S* be a subject who persists through time. For concreteness, suppose that *S* is a human subject who persists through physical time, and let all the times in the following argument be physical times indicated by some particular clock.

Suppose that *S* exists at time t_1 as indicated by some clock. (*S* also may exist before t_1 .) Suppose further that *S* persists through time until t_2 , and then splits into two essentially equal parts, which we call *T* and *U*. I will not discuss the possible means of splitting here; these details have been covered elsewhere in the literature.⁵ After the split, each of *T* and *U* exhibits behaviors that we normally regard as conclusive evidence of personhood, or at least of subjecthood. *T* and *U* continue to exist at least until time t_3 . There is no splitting of any person other than the one at t_2 .

A conventional description of this sequence of events might run as follows. Between t_1 and t_2 there is only one person, *S*. Between t_2 and t_3 there are two persons, *T* and *U*. Apparently, *T* and *U* were created when *S* divided.

If *S* really divided and *T* and *U* really were created by the splitting, then some widely discussed questions about personal identity confront us. Is either *T* or *U* the same person as *S*? Is neither one the same person as *S*? Are both *T* and *U* somehow the same person as *S*? One also can ask these questions with "person" replaced by "conscious

subject." The answers to these questions are far from obvious.

My account of the conscious subject allows us to propose answers to these questions, if we ask the questions about subjects instead of about persons. If my account is right, then the above conventional description of the splitting of a subject leaves out something crucial. That description treats the subject as though it were simply an entity which persists through time. But as we saw in Chapter 10, a subject is a temporally extended entity. A revised description of the events described above, taking into account the view of subject as history, makes questions of subject identity much more tractable.

The following argument shows us what the needed description must be like.

Consider the set V of all consciousness events involved in the splitting-subject scenario outlined above. This set includes a consciousness event x_1 which is present at t_1 . It also includes two events, x_3 and x_4 , at t_3 . At t_2 there is a fork in the series of events. One side of the fork leads to x_3 ; the other leads to x_4 . There also is a final consciousness event x_2 in the pre-split part of V — that is, in the set of all consciousness events in V which are subjectively earlier than the split. (For those who care, the proof of the preceding statement runs as follows. If there were no such final consciousness event, then there would be no final element in the segment G of V that comes before the split. Hence G would have to include an infinite series of

consciousness events approaching but never quite reaching the split. But this implies that no consciousness event in V after the split would be the successor of any of the events in G.)

One can trace an unbroken chain W of consciousness events from any consciousness events earlier than x_1 , through x_1 , through x_2 , and thence to x_3 and later consciousness events. One can trace another similar chain X from before x_1 , through x_1 , through x_2 , and thence to x_4 and later consciousness events. Each of the chains W and X is a series of consciousness events in which each member besides the very last is connected to the previous member by the continuance relation.

The part of V up to and including x_2 is like a subject history but with one difference. In Chapter 5 I pointed out that a subject history is maximal — that is, one cannot follow it to the end and then find more consciousness events that immediately succeed the one at the history's end but are not part of the history. The part of V up to and including x_2 is not maximal; one can trace through its end at x_2 and find more consciousness events which are not parts of it. The same can be said for each of the branches of V after x_2 ; both are like subject histories except that they can be extended into the subjective past. If one glues together the part of V up to x_2 with the branch of V from x_2 through x_3 , one gets a stream of consciousness which *is* maximal. But this is just W. Thus W turns out to be a subject history. By a similar argument, X is a subject history. But the part of V earlier

than x_2 is not a complete subject history, and neither of the branches of V after x_2 is a subject history. Finally, V itself is not a single stream, and therefore is not a subject history. (In the terms I used in the technical notes for Chapter 5, V is not locally linearly ordered.)

In view of these considerations, we arrive at the following description of the splitting of a subject, which is more accurate than the original description. Before t_2 there is a subject whose future includes x_3 (this subject is the history W). Before t_2 there also is a subject whose future includes x_4 (this subject is the history X). After t_2 there are the same two subjects, W and X . But W and X have an unusual feature: *they share a segment of history in common*. Every consciousness event in X 's history subjectively earlier than or at x_2 also belongs to W , and vice versa. In other words, the piece of history up to and including x_2 belongs both to W and to X . W and X do not share any parts of their histories after x_2 ; after x_2 , they act like ordinary independent subjects.

This new description embodies a possibility mentioned by Foster — namely, the possibility that dividing or merging subjects possess a common piece of history.⁶

In this new description, there is no single subject S that divides at t_2 . Instead there are two subjects before the "split" and two subjects afterward. These two subjects share the piece of subject history preceding the split. Because of this sharing of history, there *appeared* to be one subject before the split. At each moment between t_1 and t_2 , the subjects W

and X have the same spatial parts and undergo the same events. But since the two subjects are distinct as histories, they are not the same subject at any time. There were two subjects to begin with. Using standard terms of physical identity theory, one could say that W and X are constitutively identical between t_1 and t_2 though they are not numerically identical.

This description of the dividing person seems both fantastic and paradoxical. That it seems fantastic should not be cause for alarm. The very idea of dividing a person in half, body and soul, to obtain two viable persons, each possessing a piece of the original person's mind, seems fantastic to begin with. One might wonder whether such an operation can *fail* to have fantastic results. But there also are some other, more substantive objections which must be countered.

The most fundamental objection arises from the fact that before the split there exist two subjects but one series of consciousness events. On some intuitions a single consciousness should be symptomatic of a single conscious entity.⁷

My reply is this: There is not just one stream of consciousness before the split; there are two streams. A stream of consciousness is not something located entirely in the present; it cannot consist only of someone's present awareness. If a stream of consciousness is anything, it is something that goes on over time — something temporally extended. If consciousness events really are events, then the

stream is an extended event or process. Two streams of consciousness, regarded as temporally extended processes, are not identical if they differ at any time during their entire durations. In the above example the histories W and X differ in this way. Thus the numbers of subjects and of streams of consciousness in this example are equal, but there are two streams and two subjects, even before the split.

The two subjects W and X are two distinct conscious subjects who *share a common past*. This is a strange idea, but it is no more strange than the idea of a splitting subject with which we began. Actually, there is no really compelling reason why a single series of consciousness events — found by looking only at a short segment of time and ignoring the past and future — has to indicate the presence of only one subject. Since a stream of consciousness is something that stretches out over time, streams of consciousness which differ at any time should be regarded as different streams of consciousness, even if they sometimes share the same events. But a subject is just a stream of consciousness — so two different subjects may share consciousness events in common and still remain distinct, provided they do not share *all* of their consciousness events.

One can imagine a universe containing exactly two subjects having exactly similar, indistinguishable experiences. It is not intuitively clear whether there really are two subjects in this case, or whether there is only one. (This example is analogous to Max Black's example — proposed for a different purpose — of two spheres, identical

in their qualities, in an otherwise empty world. Black asked, in effect, whether there really is any distinction between these spheres.⁸) But if at any moment the two subjects differ in their experiences, then they cannot be regarded as the same subject. Although there are few intuitions suggesting that conscious beings can share parts of their histories, there are no really strong intuitions suggesting that this is impossible, especially in the intrinsically counterintuitive case of the splitting subject.

Someone might object that two subjects which are indiscernible at any time have to be identical at that time. (A traditional philosophical principle, the law of identity of indiscernibles⁹, might lead one to think that way.) My reply to this is that the two subjects W and X are not indiscernible even before the split. They cannot be said to live the same life at any time. This is because they always have different *futures*. Before the split, W and X may be indiscernible with respect to some properties, but not with respect to all properties which refer to the future. For example, if after the separation X goes on to become a professional hit man and W goes on to graduate school instead, then at any time before the split it is the case that X *will be* a professional criminal and that W *will not*. In lay people's terms, X is a budding hit man and W is a budding professor. Such future-referring present properties can *individuate* (establish separate identities for) subjects. To deny that they can is to deny that subjects include their entire histories; this contradicts my account of subjects. But aside from my account, such a denial seems to run counter to intuitions

which suggest that a person's future path is an important feature of that person. To be "going somewhere," headed for a particular destiny or outcome, is a property a person may have now. Intuitively, it sometimes seems that persons who are very similar in their present attributes, yet doomed to very different fates, are different by virtue of these non-present differences alone. It is not obvious that subjects who have largely the same present properties at a particular time, but who differ with respect to some important property involving the future, are the same subject now. (Note that none of what I have just said depends upon the existence or nonexistence of free will, or upon the truth or falsity of determinism.)

The above discussion points out a reason why we should not say that W and X "live the same life" before the split. At all times before the split, W and X have different futures; hence their lives, which include their futures, are different. (Ask yourself thoughtfully whether a future hit man and a future professor could be said to be living exactly the same life now, even if their lives are the same in all *present* particulars.) If one maintains that "living the same life" means having the same present experiences, then W and X live the same life before the split, but then it is not strange that they can do so and still be distinct.

Could I Be Two People?

A more serious potential problem with this view of the

splitting subject arises from the possibility that the two subjects before the split may believe that they are one.¹⁰ Since the subjects before the split share the same consciousness events, any thought or other experience had by one during that time will be had by the other as well. For example, the two subjects may simultaneously decide, on the basis of easily available evidence, that they are *one* subject. In this situation, they may simultaneously call themselves "I." So how do I know that I am now one subject, and am not really two subjects who have not yet divided? If a person living now divides sometime in the future, can he justifiably conclude that he always had been two persons after all?

A related problem stems from the observation that the splitting of a subject can appear to influence the past. If it is possible that I might split in the future, then my being one subject or two subjects *now* appears to depend upon whether or not a certain event will occur in the future.

These difficulties vanish when one looks at them more carefully. I will address them in turn.

First, I can be sure that I am one subject and not two because the fact athat I am one subject follows from what I know about my immediate past consciousness events. (In Chapter 4 I showed that some such knowledge is trustworthy.) I am now in the midst of one consciousness event. I am experiencing in continuance one consciousness event. From this knowledge I can infer that I am not now identical to two distinct subjects. Of course, if by "identical" I mean "having the same present parts now," then this

conclusion may be wrong. I cannot know for certain that another subject is not sharing my consciousness events; to know that, I would have to know that I will not undergo a brain split some day. What I can know by reflection on direct experience is that what I now call "I" is a single subject — that when I utter "I am *this* subject," the statement uttered refers to a single stream of consciousness. From this it follows that I am not a cluster of subjects, and that my usage of "I" is not ambiguous between two or more subjects.

This intuitive insight remains true even if there is another subject having exactly the same consciousness events that I am having. If there is such a subject, that subject will not be the subject which *I* just called "I." Of course, *he* will say "I," or "I am this subject," at the same moment when I do; but when *he* says it, he is referring to *himself* — a numerically distinct piece of subject history. Since "I" is a word which always refers to the speaker, no one lies.

If you are going to split, you already are one of the subjects who will exist after the split. If you are the incoming "subject" S in the above example of splitting, then in reality you now are either W or X. Of course, you may not know now which one you are. But after the split, you will begin experiencing either the life of W or the life of X, and not both. If we fall back on the conventional and erroneous description in terms of S, T, and U, we can describe the outcome this way: Before the split you are S; after the split, you will become either T or U, but not both.

If you are to divide in the future (more correctly, if a process which you ordinarily would call "division of myself"

is to happen in the future), then you will be one of the subjects which will result from that process, but you will not be both of those subjects. You already are the subject you will be after the split; you will not become a different subject after the split. In other words, you, as a conscious subject, will not really divide. The other subject who will result from the split also exists now and shares your present consciousness events.

As for the possibility of changing the past, note that *your* being one subject or two subjects now does not depend on the happening or nonhappening of a future event. You are one subject, and no future event can alter this fact. The happening of the future event may bear on the presence now of other subjects who share your present consciousness. But this does not involve any real change in the past, for the consciousness events happening to you now are not themselves changed in any way. The later split may result in an earlier consciousness event's belonging to two different histories, but that does not constitute a change in the past. Also, since histories exist tenselessly, the number of subjects which exist now is the same number which exist after the split.

If one still wants to say that the split "creates" a new subject, one can say that only in the following sense: if the split in a stream of consciousness did not occur, there would only be one subject in that stream.

Despite its bizarre, counter-commonsensical look, this account of the dividing subject apparently involves no inconsistency and does not really violate our most central

intuitions about subjects. It appears to clear up entirely the problem of subject identity in the case of a dividing subject. In conventional language, the upshot of the solution is this: If a subject divides, that subject will end up as one of the subjects resulting from the division.

Splitting Subjects with Psychological Compartmentation

The above account of the dividing subject also suggests a way in which a subject might predict in advance which product of a split he or she will become. This way works only for certain kinds of subjects: those whose mental processes are divided into two or more subprocesses, each of which is partially independent of the other. I will argue that human conscious subjects are subjects of this sort on account of the functional division between the two hemispheres of their brains. This section is even more technical than the preceding one, but those interested in the split-brain problem should not skip it.¹¹

Let us return to the splitting-subject example detailed above. When S divides at t_2 , any subject having S's history up to that time finds that he is either W or X. Which one he is depends upon whether he was W or X to begin with. Now suppose that all consciousness events in the example really are events (as they are for human conscious subjects), and that each consciousness event which W or X undergoes before t_2 is a composite of two events, each of which is not a

consciousness event. Specifically, suppose that the instance of seeming which is the consciousness event involves *two different kinds* of appearance; the facts which are the case for the event can be partitioned into two classes which are different in some significant way, and the subject's being aware of the facts in one of these classes is an event. It is not hard to find consciousness events which probably are of this sort. For example, some of our consciousness events involve acts of thought accompanied by emotional states. In such a case the occurrence of the act of thought is an event and the occurrence of the emotion is a distinct event; the consciousness event consists in the occurrence of both. (Of course, these other two events need not be *spatial* parts of the consciousness event, although if they are neural events, then they might be such parts.) Another example: a consciousness event might involve the seeing and hearing of the same external event.

Suppose furthermore that these partial consciousness events form two chains in the following manner.

Let y and z be successive consciousness events in the subject's history, with y continued in z . Let y_a and y_b be the partial consciousness events comprising y , and likewise z_a and z_b for z . During z it seems that y_a just occurred in the immediate past — that is, that some of the subjective facts of y , which also are involved in y_a , were the case in the immediate past. Similarly, it seems during z that y_b just occurred. The fact that y_a just occurred is among the facts whose being the case for z defines z_a . But it is *not* among

the facts whose being the case for z defines z_b . In other words: the event y_a is in the subjective "realm" associated with the event z_a , but is not in the subjective "realm" associated with the event z_b . Similarly, the event y_b is part of what is experienced in the event z_b , but not in the event z_a .

We can suppose that each of the consciousness events prior to t_2 is divided into an "a-part" and a "b-part" in this way, and that the "a-parts" of successive consciousness events in the chain are linked together in the way described in the preceding paragraph. We assume the same for the "b-parts." Intuitively, this means that the subjective life of the subject is divided into two streams of experience which the subject does not witness together. Such division of experience is a special case of the psychological compartmentation discussed in Chapter 11.

Just for convenience, we will call these two chains the *right hemisphere* and the *left hemisphere* instead of the chain of "a-parts" and the chain of "b-parts." (These terms should not be taken too literally; the subjects involved may be Martians who lack brain hemispheres as we know them. Later, my real motive for using these terms will become manifest.)

Recall that x_2 is the final consciousness event in the part of W 's and X 's history before the split. Let x_a and x_b be the a- and b-parts comprising x_2 , with x_a in the right hemisphere and x_b in the left hemisphere. Suppose that the split occurs

in such a way that the first consciousness event (say y_W) in the branch belonging to W alone involves continuance of x_a , but not of x_b . (That is, y_W holds x_2 in continuance and is a successor of x_2 , but it does not seem at y_W that *all* subjective facts true at x_2 were true; it only seems that way about *some* of those facts, specifically some facts belonging to the "part" x_a .) Suppose similarly that the first consciousness event (say y_X) in the branch belonging to X involves continuance of x_b but not of x_a . Both y_X and y_W continue x_2 and thus are successors of x_2 , but each one continues only some of the contents of x_2 — the part associated with one of the two partial events associated with x_2 .

The hypotheses of this thought experiment imply that before the split, the facts which are the case for W and X belong to two "compartments," in the psychological sense of that word. We may speak of a given subjective fact as belonging to the a-compartment or to the b-compartment. By this, we do not mean that there are actual entities called "compartments," although if one must, one can regard the two compartments as properties, classes, or predicates of subjective facts. We can say that for y_W , contents of the a-compartment, but not those of the b-compartment, are continued. (By this we mean only that the event y_W involves the awareness, in continuance, of some facts in the a-compartment in the manner described above, but not of any facts in the b-compartment.) Similarly, y_X involves the continuance of the contents of the b-compartment but not

those of the a-compartment. At the moment after the split, W is the subject whose immediate past contains contents from the a-compartment, but none from the b-compartment. Analogous statements hold for X.

Now we have a way to trace the careers of the subjects W and X before the split. Suppose it is I who split. If immediately before the split, I *know* that I am experiencing the contents of the a-compartment only, then my present conscious contents are linked, in the way described above, to the contents of W. If immediately before the split I am knowingly experiencing only the contents of the b-compartment, then my present consciousness is a predecessor, in the same way, of the consciousness of X. Hence if before the split I am aware of experiencing the contents of the a-compartment, after the split I will be W. If before the split I am aware of experiencing the contents the b-compartment, then I am to be X.

The alternatives to these outcomes are less plausible than the outcomes just proposed. If I experience the a-compartment but I turn out to be X, then X's history after the split begins from x_a , which, by hypothesis, X never experiences in continuance after x_a . Similarly for W and b.

Splitting Subjects Again, in Less Detail

The argument of the last section has the following consequence. Suppose that the contents of a subject's consciousness fall into two groups — that is, two different,

mutually exclusive kinds of subjective facts. Now suppose that that subject splits in such a way that each final subject experiences, in continuance, only one of those groups of facts. Under these conditions, each of the two groups of subjective facts will belong to the past of only one of the final subjects. If just before the split, a subject is aware of having only one of the two kinds of experience, then that subject will become the final subject which has that kind of experience in its past.

The scenario of the last section is, of course, designed to resemble the division of a human being into two halves, each controlled by a single brain hemisphere. To make the resemblance clear I will list some well-known background information from neurophysiology.¹² The right and left hemispheres of the human brain perform different functions in mental life. Our mental processes can be classified, roughly and nonexhaustively, into two sorts. Some processes, such as spatial perception, are associated primarily with the functioning of the right cerebral hemisphere. Other processes, such as language production, are connected similarly with the left hemisphere. The division of mental processes along hemispheric lines is neither exhaustive nor anywhere near exclusive, but its approximate correctness seems to be well-established.

The two sets of subjective facts associated with these two kinds of processes correspond to the two compartments in my argument. If one grants this correspondence, then the mental processes associated with the functioning of the two hemispheres are the processes which I called "hemispheres"

in the argument. (Hereafter I will call them *hemispheric processes*.) One's having of an experience via either left or right hemispheric processes constitutes an event. These are the events which belong to the left and right hemispheric processes. During any consciousness event, two events occur; experience via the right hemisphere constitutes one event, and experience via the left hemisphere constitutes another event. The simultaneous occurrence of these two events is sufficient for the occurrence of a consciousness event in the subject. Of course, this correspondence between the division of a real brain and the hypothetical division described in the last section is a gross oversimplification, since the hemispheres are not tight compartments — they do not even come close to complete separation. But subjective facts which are associated with the joint functioning of the hemispheres can be assigned to both compartments without changing anything essential in my argument.

Note that in the above arguments about hemispheres, the two events which comprise a consciousness event are not themselves consciousness events of the subject. A single instance of seeming in the subject's experience may involve subjective facts originating with both hemispheres. One might think of a consciousness event in a normal human mind as constituted of two events (or groups of events), neither of which is itself a consciousness event and each of which is an event associated mainly with the activity of one brain hemisphere.

Human beings differ from the double subject of the

example in another way as well. In addition to having leaky and overlapping hemispheric processes, humans are known to exhibit so-called *hemispheric dominance*.¹³ John C. Eccles has noted an important aspect of this dominance: we experience the contents associated with one hemisphere (usually the left one) more readily than those associated with the other hemisphere. Eccles has suggested that neural processes which are conscious occur in the dominant hemisphere, while processes in the opposite hemisphere are in themselves unconscious.¹⁴ This suggests that if I am a human conscious subject, my dominant hemisphere's compartment is the one which I normally experience in continuance. We may conclude that if I am split in such a way that one hemisphere remains in control of each of the resulting hemi-beings, my consciousness will go with that product of division which includes my dominant hemisphere.

Eccles has suggested, on the basis of clinical evidence, a similar conclusion about the results of dividing the brain. From clinical evidence, he concludes that in cases of split brain or of the loss of one hemisphere, the subject's conscious functions typically go with the dominant hemisphere.¹⁵ This supports the suggestion which I made at the end of the preceding paragraph.

This argument may be extended to beings which lack brain hemispheres like ours. If such a subject has two or more "hemispheric" processes, one of which is more readily accessible to the subject, and the division isolates the more

accessible process in one of the products of division, then the product having that process will be the subject.

Commissurotomy and Subject Splitting

Cases of human commissurotomy (split-brain surgery) differ in significant respects from genuine subject splitting. In commissurotomy, the processes associated with the two hemispheres still can interact with each other causally; they clearly influence each other in ways which presumably alter their subjective contents.¹⁶ A split-brain patient really is no more divided than are some of the psychologically compartmented ordinary subjects whom I discussed in the previous chapter.¹⁷ Hence it is rash to regard the processes in the two separated hemispheres as two different series of consciousness events. Split-brain patients generally seem to be directly conscious of perceptions associated with their dominant hemispheres, but not of processes in the opposite hemisphere.¹⁸ Yet even if this were not the case, it is unlikely that two streams of consciousness would exist in the history of the patient; there still would be one subject. The worst that has happened is a very dramatic *compartmentation of content* of the same sort that I discussed in Chapter 11.

If the two hemispheres were *completely* severed so that each one could function without even indirect communication with the other, then two subjects might exist. On my model of the splitting subject, the patient we

knew before the operation would continue as the half of him/herself controlled by the dominant hemisphere. The weaker hemisphere would continue to live as a separate subject, with (for reasons discussed above) different capabilities.

Merging Subjects

Puzzle cases in which subjects merge can be treated along the same lines as those in which subjects divide. Such merging-subject cases have been discussed in the literature.¹⁹ The following scenario is typical. At time t_1 there are two persons A and B. At a time t_2 later than t_1 , A and B fuse (not necessarily instantaneously). At times later than t_2 there is only one person C, whose history begins with the fusion event at t_2 .

We can redescribe this chain of events in terms of temporally extended subjects, much as we did for the case of the dividing subject. The resulting description runs like this. There are two subjects, D and E. The parts of these subjects' histories earlier than t_2 are distinct; D and E do not contain those parts of their histories in common. The parts of the histories of D and E at and after t_2 belong both to D and to E. In other words, at and after t_2 , D and E share a common future.

In this scenario, the "person" resulting from fusion actually is two subjects, just as is the "person" who has not yet divided in the previous scenario. When two subjects

fuse, the result is two subjects sharing a future. Neither of the original subjects ceases to exist, and no new subject is created. The being resulting from such a fusion is much like the initial being in the dividing-subject case. Both of these beings are composites of two subjects which behave like one subject but nevertheless are not identical to one another.

The fact that neither subject ceases to exist suggests that it may be logically possible to undo the fusion. Since neither subject is lost, there is no reason to believe that the subjects could not later separate. This possibility directs our attention toward a more confusing (and also previously studied) puzzle case. What happens when two subjects merge and the being resulting from their merger subsequently divides?²⁰ Can either of the incoming subjects be identified with either of the outgoing ones?

An analysis similar to the ones above shows that in this puzzle case there are *four* subjects, and that no subject is created or destroyed. The following scenario illustrates this fact.

Suppose that at t_1 there are persons F and G. At a later time t_2 , F and G undergo fusion to form a being H. Later still, at t_3 , H divides to produce two new conscious beings, I and J. Can we say whether F or G is H, whether F or G is I or J or neither, and whether H is I or J or neither?

Here is the same scenario redescribed in terms of temporally extended subjects. There are four subjects. One of these subjects has a history consisting of what we call the history of F before the merger, followed by the history of H,

followed by the history of I after the split. In an obvious notation we may call this subject FHI. Another subject has in its history the history of F before the merger, then the history of H, then the history of J after the split. We call this subject FHJ. By similar means we trace the subject histories GHI and GHJ. Each of these four histories counts as a subject. There are four subjects before the merger, four subjects between merger and split, and four subjects after the split.

We see now that the the initial "person" F actually is two subjects leading one conscious life. Those two subjects are the ones which share the history of F as parts of their own histories, namely FHI and FHJ. The beings at G, I and J also are dual subjects of this sort; each of them is a pair of merged subjects. The being at H is a composite of four subjects sharing a common piece of history: it is a phase in the life of each of FHI, FHJ, GHI, and GHJ. The objections which arise in the splitting case also arise (with appropriate changes) when more than two subjects share a piece of subject history. These objections do not seem any more troubling in the four-subject case than in the two-subject case. The arguments which defuse them in the splitting case can be adapted to accommodate more than two subjects.

We also can imagine a case in which two subjects merge and redivide without any blurring of their original identities. This can happen if the initial, final, and middle segments of subject history have dominant hemispheric processes (with "hemispheric processes" defined as in the splitting case). Suppose that the dominant hemispheric processes of F and

G are continuous with the two hemispheric processes of H, in the sense that the initial event in each hemispheric process in H subjectively succeeds either the final event in the dominant hemispheric process in F or the same thing in G. If later on, the hemispheres of H part company to yield two separate beings, then one can make a case for the conclusion that one gets back the same subjects (also known as F and G) that went in. I leave the details of this to the reader since they are quite similar to the splitting of the subject with hemispheric dominance. There now are two subjects; one of the two is simply F and the other one of the two is simply G.

In Conclusion

The accounts of division and merger of subjects which I have presented here have a fantastic air about them. I would argue that from the viewpoint of an adequate understanding of the conscious subject, they are not fantastic at all. At least they are no more fantastic than the already bizarre ideas of splitting one person into two and of merging two persons into one. Such happenings, by their very nature, stretch our everyday concepts of personal identity to their outermost limits. Ordinarily one does not even conceive of the possibility that two or more distinct subjects could have exactly the same experiences during an interval of time, or that what happens to a body in the future determines how many subjects it houses now. But any uneasiness should disappear when the reader recalls that I have identified

subjects with their *histories*, and re-reads my arguments with this fact firmly in mind. It is not mysterious for two *histories* to have some events in common; certainly World War II is an event in the history of England and also in that of France. Nor is it strange when a single series of events, by virtue of its relationship to other events, is a segment of two distinct histories. The history of the United States coincided with that of the thirteen original colonies only for a circumscribed period.

If one forgets that subjects are histories, then my description of what happens when subjects split will seem to imply that before a human being splits, two minds always exist in the same body — or something like that. Actually, I have made no such incredible claim. The most I would claim is that before a human being splits, two *personal histories* are exemplified by one *consciousness*.

Everyday thought, and some philosophical views as well, regard a subject as something without temporal parts.²¹ On my account, subjects do have temporal parts; they are coextensive with their histories. My assignment to subjects of properties typical of histories is what lends my accounts of splitting and merging subjects their air of unfamiliarity. But the absence of this same feature in other accounts of personal identity makes the problems of splitting and merging subjects seem far more difficult than they really are. On my account, a conscious subject is something which exists entirely at each moment in its history but also is temporally extended. The compatibility of these two features of subjects is a fortunate consequence of my

From Brain to Cosmos

nonstandard view of persistence through time. If one accepts the identification of subjects with subject histories, then the above stories about what happens when subjects fuse and divide lose much of their oddness — and the puzzle cases about splitting and merging subjects become far more tractable.

Chapter 13

Mind and Matter

In the past several chapters I have discussed a variety of questions about consciousness and time. The conclusions at which I arrived provide partial answers to the question which I posed in Chapter 1: "What can we learn about the nature of reality by deducing the consequences of facts about how things seem to conscious beings?" Beginning with the logical properties of subjective fact and with descriptions of how certain situations appear to observers, I have been able to develop a theory of conscious subjects and some accounts of various features of time. According to these accounts, many of our commonsense beliefs about the world around us are true. In particular, we live in a real world inhabited by conscious subjects which persist through some form of time, and which have experiences of physical objects which at least seem to exist and to persist through time. (Those who regard these findings as obviously true, and who regard the last sentence as troubling on these grounds, should re-read Chapter 1. The important point is not *that* we established

these obvious conclusions, but *how* we established them.)

So far I have sidestepped most of the "big" problems of traditional philosophy. In this chapter I will confront one of these problems: the question of the relationship between the mind and external reality. Traditionally, answers to this question have been of two sorts: *idealistic* ones, according to which mind is the underlying reality which somehow gives everything else existence, and *realistic* ones, according to which other things besides minds are truly real. In this chapter I will present a new solution to this problem of the link between mind and reality. This solution is the centerpiece of the new view of reality which I promised in Chapter 1.

The Notion of Objective Truth

Before I start to examine the relationship between mind and reality, I want to say a few words about the philosophical problem of truth.¹ The idea of truth will come up repeatedly in the later part of this chapter, and I want to prevent certain misunderstandings before they start.

The problem of the relationship between mind and reality is intimately connected with the problem of determining the truth conditions for statements about real objects — that is, the conditions under which those statements are true or false. Up until now, we have dealt largely with statements about how things seem or about instances of seeming. The truth conditions for such statements have not concerned us,

since we know how to tell whether such propositions are true or false. But if we want to ask ourselves what role mind plays in the makeup of objective reality, we must be able to say what, if anything, makes a statement about objectively real entities (not merely apparent ones) true or false. That is, we must know what the truth conditions for such statements are like.

Everyone has an intuitive feel for what "true" means. However, different people have different views about how truth may be arrived at, and about what makes a statement true or false. The existence of differences in prephilosophical usages of "true" and "truth" has suggested to philosophers that those words are somewhat ambiguous.² The following scenario shows how disagreements can arise from ambiguity of this sort.

Bill, a person who lives by sense experience, correctly asserts "The cat is black." When asked to explain why this claim is true (not merely to show that it is true), Bill asserts that it is true because the cat really *is* black, and that this is all that needs to be said. Then Bob, who is a religious believer, claims that "God exists" is true. When asked to explain why this claim is true (not merely why he believes it), Bob replies that belief in God is indispensable for making sense out of life — that without this belief, one will have great difficulty understanding the world or finding meaning in it. Bob might further state that the truth of the claim that God exists is not a question of "just facts."³ Upon hearing this opinion, Bill might reply that "God exists" cannot be true, and that it cannot be true precisely *because*

God's existence is not just a matter of "facts."

In this discussion, Bill and Bob are disagreeing about theology, but they are not disagreeing over theological points alone. Their disagreement stems from diverging insights about why statements are true or false — about what makes a statement true or false. Bill appears to think that for a statement to be true it is necessary and sufficient that it corresponds to the facts. Bob appears to think that this condition is not necessary; instead, he thinks it sufficient that the statement makes sense out of, or fits in well with, other truths.

The intuitions about truth which seem to be motivating Bill and Bob are reminiscent of two of the standard philosophical theories of truth.⁴ Bill's intuitions resemble an informal version of the correspondence theory of truth, while Bob's intuitions remind one of the coherence theory. In real life, one finds people with intuitive views about truth which remind one, in greater or lesser degree, of these theories. After noticing these differences of intuitions, one cannot help but wonder whether the intractability of the philosophical problem of truth might arise, at least in part, from the existence of these varying intuitions about the nature of truth.

Many fields of human endeavor, ranging from theoretical physics to law, require their practitioners to think about what is true. The particular role which truth plays in a given discipline often brings to mind one of the philosophical theories of truth. The idea of evidence used in experimental science seems natural if one accepts the correspondence

theory of truth. The methods which pure mathematicians use to determine the truth may look more natural on a coherence theory. The sort of truth with which engineers must cope every day looks much like the picture of truth painted by pragmatists. And formal logic is easiest to interpret when truth is given by Tarski's semantic definition.⁵ These differences in the roles which truth plays in different fields suggest that perhaps there is no unique answer to the question "What makes a statement true?"

Before continuing this discussion, I should say a word about my terminology. By "statement" I mean a linguistic item which is true or false. I will not specify the nature of these items (for example, whether they are tokens or types), because I do not need to so for what follows. I use "statement" instead of "sentence" or "formula" because I want to avoid the overtones which the latter two words have acquired from their association with formal logic. A reader who prefers to think of truth as a property of propositions could replace "statement" with "proposition" and still make use of my arguments.

Continuing the argument: Perhaps there is more than one general kind of circumstances that can make a statement true. Perhaps some statements (say, those of pure mathematics) are made true by coherence of some sort, while others (those of physical science, say) are made true by correspondence with fact. It is conceivable that none of the philosophical theories of truth exhausts the informal notion of truth, but that this notion still is based upon a perfectly sound intuition which everyone who thinks can

have. A precise definition of truth might best be regarded as a definition of truth for one class of statements.

Note that the argument of the last paragraph *does not* endorse the relativist opinion that truth is merely relative or is a matter of viewpoint. Even if no single theory captured the entire notion of truth, there still might be a single, objective notion of truth. Such a notion would remain tenable because the different theories of truth never would deliver different verdicts on the truth of a single statement. If two supposed theories of truth did disagree in this way, no one could regard both of them as adequate theories of truth, for if either one were right, one could use it to show that the other theory is inadequate. The applicability of different theories of truth for different kinds of statements would not imply that there is more than one concept of truth, or that truth is relative. All it would imply is that truth is too rich a notion to be captured entirely by one theory or to be approached solely through one method. It would mean that no rigorous theory of truth is powerful enough to exhaust the content of the informal, intuitive notion of truth, which is the complete notion. Truth might be, in this sense, *unbounded*.

A situation like the one described in the last paragraph already exists in the field of axiomatic set theory. Set theorists know that in the standard formulation of set theory (ZF) there is no rigorous, formal characterization of truth in a model which assigns all formulas their correct truth values.⁶ In other words, the notion of truth used in ZF set theory cannot be exhausted by a single rigorous definition of

truth statable in the language of that theory. Yet this does not lead set theorists to regard truth in ZF as a mere matter of opinion, or as entirely viewpoint-dependent or relative. Set theorists are right in not thinking that way; the mere fact that truth in ZF cannot be formalized all at once in one definition does not justify regarding that notion as anything other than objective. If ZF is a consistent theory, then the different "theories of truth" (or formalizations of the notion of truth) available in ZF never will disagree on the truth of any formula. (If two of them did assign the same formula two different truth values, then they could not both be genuine truth definitions, for if one definition were right it would allow us to prove that the other was wrong.) The fact that the notion of truth in ZF cannot be captured by a single truth definition within ZF does not mean that there is more than one conflicting notion of truth in ZF, or that there are many divergent truths about what is a theorem of ZF, or that there is no objective notion of truth in ZF. It simply means that the notion of truth in ZF is too rich to capture in a single definition.

Another relevant example from formal logic has to do with the decidability of formal theories. Many formal logical theories are undecidable — that is, one cannot always tell whether a given wff (or generalized sentence of symbolic logic) is deducible from the theory, or at least one cannot always tell by rigorous, mechanical means.⁷ Different kinds of proofs are needed for different theorems; there is no single rigorous test which will tell us whether a given wff is a theorem. Yet it would be silly to infer from

this that there really are no theorems in such systems, or that the concept of theoremhood really is a set of different, irreconcilable ideas.

The conceptions of truth embodied in the various philosophical theories of truth are not logically equivalent to one another. If they are jointly right in the way I have suggested, then they are applicable to disjoint classes of statements, or else they must agree on any statement to which two or more of them are applicable. If two such conceptions assigned such a statement different truth values, then both conceptions could not be right, at least not in the absence of some amendments or caveats. Each of the theories of truth would capture certain features which common sense attributes to truth, but would overlook some other features of the same kind. The above example of Bill and Bob suggests how one can arrive at different theories by trying to rigorize the informal notion of truth. The correspondence and coherence theories describe what Bill and Bob might have meant by "truth" had they reflected carefully on their opinions and spoken more precisely.

This view that truth is *single and objective but multifaceted* is not essential to what follows. I am not going to defend this view any more than I already have. My point in discussing this view of truth is to show that the search for a single correct theory of truth may not be the best way to approach the problem of truth. Even more misguided are the relativistic attempts to debunk the notion of truth by showing that there are many different kinds of truth. Acceptance of different theories of truth for different kinds

of statements does not push us toward relativism *of any sort whatsoever*. It also does not immediately rule out relativism, although it undermines one particular line of argument for relativism. (Some relativists might misinterpret my views as a kind of relativism — but that is their problem, not mine.)

The rest of this chapter does not presuppose any particular theory of truth.

From Appearance to Knowledge⁸

To know anything, a conscious subject must have experiences of appropriate sorts. Information which never enters the consciousness of a subject cannot become part of what that subject genuinely knows. I am speaking here only of *discursive* knowledge — the kind of knowledge which philosophers typically profess to seek. There also are non-discursive forms of knowledge, such as the knowledge of how to ride a bicycle. (Some people would not regard this last form of knowledge as genuine knowledge at all, but their beliefs on this point do not prevent them from knowing how to ride a bicycle.)

If you find the answer to an arithmetic problem with the help of a nonconscious procedure like a mechanized calculation, then you do not know that answer until you become conscious of it. If a nonconscious machine carried out the same calculation in deep space, unexperienced by any subject, you could not be said to know anything as a

result of those operations. You can know only what affects your world of appearances — what has an impact upon the way things seem to you. It is only through conscious experience — through a way things seem — that a subject can know anything at all.

Suppose that the truth of a statement P does not have any consequences for the way things seem to conscious subjects. That is, suppose that every experience had by every conscious subject will be exactly the same whether P is true or false. Then there is no way for a conscious subject to know that P is true. No one ever will find any evidence for the truth of P.

The possibility that a subject would "just know" that P, without having to refer to any evidence, might seem to provide a loophole in what I have just said. But even if intuitive knowledge of this unsupported kind were possible (and I doubt that it is), still it would not weaken the argument in the last paragraph. Suppose that some subject "just knows" that P, and this knowing that P is brought about somehow by the fact that P. (Perhaps biological evolution forced this subject to believe that P; this could happen if P is a persistent fact about the natural environment and natural selection favored genes conferring this knowledge.) Then things seem different, to the subject, from the way they would seem if P were not the case; hence the subject's experience is not the same as it would be if P were false. In this case, the subject's knowing that P would not contradict what I said in the last paragraph. Alternatively, suppose that someone "just knows" that P, and this knowing that P is *not*

causally linked, or correlated in any other way, to the fact that P. Then if P were false, the experience of "knowing" still might occur. Such an instance of "knowing" is not a genuine instance of knowledge, but is merely an instance of belief backed by some sort of psychological compulsion.

If the truth of a statement (or proposition) P does not follow from some set of facts about how things seem to a subject, then regardless of how things seem to that subject, that subject cannot absolutely rule out the possibility that P is false. Therefore, a subject cannot know for certain that P is true unless P follows from facts about how things seem to that subject. Of course, the subject might be able to begin with facts about how things seem and infer in a less-than-certain way (non-deductively) that P is true. But this cannot happen unless P follows from facts about how things seem to the subject, by way of whatever kind of inference the subject is using. For example, if the subject is using inductive reasoning, then it must follow, from facts about how things seem to the subject, that the truth of P is probable. If the subject is using some kind of intuitive method, and if we assume, for the sake of argument, that this method can be trusted, then facts about how things seem to the subject must somehow render P intuitively acceptable or convincing. (Perhaps P simply *seems* true.) If P cannot be justified in any way at all beginning from facts about how things seem, then it could be the case that things seem exactly the way that they do and yet P is false. In this case, there are no real grounds for belief that P.

The above arguments show that a subject can know that a

statement is true only if that statement can be justified, in some way or other, beginning from facts about how things seem. (I will not argue here for the superiority of any particular method of justification; my point is independent of such controversies as empiricism vs. rationalism, or empiricism and rationalism vs. mysticism.) If a subject can know with certainty that a statement is true, then it must be possible for that subject to infer that statement with certainty from facts about how things seem. If a subject can know that a statement is probable, then it must be possible for that subject to infer from such facts that the statement is probable. If a subject can know on some other grounds that a statement is credible, then it must be possible to infer, from facts about how things seem, that the statement is credible on those grounds. If the truth of a statement has no bearing whatsoever on how things seem to a subject, then that subject does not have justification for supposing that that statement is true.

All this shows that facts about what is the case for subjects can be used to evaluate the truth value of any statement, *to the extent to which that truth value can be known.*

It follows that if we had a complete description of how things seem to all subjects, then we would be able to determine the truth value of any statement whose truth value is knowable, and we would be able to do this with the highest degree of certainty possible. (Of course, such a description is a practical impossibility, although in principle we can get as big a finite fragment of it as we need.) We can

determine the truth value of any statement, to the extent that it is knowable, from facts about what is the case for various subjects. But a fact which is the case for a subject is the case for that subject only at a particular consciousness event. Therefore, the truth value of any statement or statement can in principle be determined, to the extent that it is knowable, from enough data about what is the case for consciousness events.

The preceding remark is not meant to suggest that all the required data about "what is the case for what" can be expressed in a particular language. It is safe to suppose that no language can express *all* subjective facts. (A language which could do this would have to be able to express every subtle shade of how every subject feels at every time.) Nevertheless, if we want to talk about a particular subjective fact, we always can manage to do so by extending our language a bit. If we want to express how things seem to someone by using a sentence of the form "For x, it is the case that P," we need only invent a new word or other symbol to express the fact which is the case for x. People often make up or adopt new words or phrases to express things which previously were inexpressible for them. Children do this all the time as they learn language. For anything which seems to be the case, we can invent a sentence to express that "something". If worse comes to worst, we can simply invent a new sentence letter or mathematical symbol, and use it to express the previously inexpressible fact that _____.

All of the preceding arguments lead up to the following

conclusion.

Let P be a statement. Suppose that it is logically possible that some subject knows the truth value of P. Then the truth value of P can be determined, to the extent to which it can be known, from facts about what is the case for consciousness events.

I will call this conclusion the *experiential principle*.

The Experiential Principle: What It Is Not

The experiential principle has some interesting consequences. Before exploring these consequences, I wish to mention some consequences that the principle does not have. Each of the views discussed here bears some superficial resemblance to the experiential principle, but should not be confused with it. All of them differ from it in significant ways.

First, the experiential principle does not imply that any form of *empiricism* is correct. This is because it does not imply that sense experiences are the only experiences which could bear on the truth of a statement. The subjective facts which the principle requires could be those associated with sense experience, but for all we know they could just as well be those involved in thought, emotion, intuition, mystical experience, or Platonic recollection. (I am not arguing here that all of these sources of knowledge are equally

meritorious; I am simply claiming that the experiential principle alone cannot be used to rule out any of them.) The experiential principle does not specify what *kinds* of subjective experience can yield knowledge. Hence it neither supports nor conflicts with empiricism. (Note particularly that it does not imply positivism or verificationism. Far from it.)

Second, the experiential principle is not a *foundationalistic* thesis. It does not imply that one can begin with facts about what is the case for whom and build up all knowledge from them. The procedure which leads from subjective facts to the truth value of a given statement need not be a deductive one or even a rational one. For all we know, it could be a scientific procedure involving intuitive leaps and guesswork, or even an intuitive method (Bergsonian, phenomenological, or other) by which one plumbs experiences for their inner meanings. I am not arguing for or against the merits of any of these methods; I am simply pointing out that the experiential principle does not rule out the use of such methods as paths to knowledge. The open possibility that some of our knowledge may be available *only* through such methods may not be compatible with foundationalism as we know it.

The experiential principle tells us that if it is possible in principle to know the truth value of a statement, then it is possible in principle to determine that truth value from facts about how things seem to subjects. It does not specify how this determination is to be done.

Third, the experiential principle does not imply that one

can know only about one's own experience. The principle says that any fact that one can know discursively, one can in principle derive from facts about conscious experience. This may appear to harbor the beginnings of solipsism, but the experiential principle does not imply that one knows only about the contents of one's own consciousness. It allows one to know about external things *through* the contents of consciousness, and that means the consciousness of other subjects (indirectly) as well as that of oneself. For example, you might be able to ascertain the existence of an object of which you are not aware, by examining the reactions of other persons to the object. The resulting inference may indirectly utilize other people's subjective facts (which, as I pointed out in Chapter 6, may be more accessible to you than you think). Of course, your knowledge of other peoples' subjective facts is obtained with the help of your own subjective facts.

Fourth, the experiential principle does not imply or even suggest that only perceptible objects exist. The principle is compatible with the possibility of knowledge about imperceptible things. The experiential principle tells us that any such knowledge would have to be derived, deduced, induced, intuited, or obtained in some other way from facts about what is the case for consciousness events. For example, invoking an object as part of a theory which explains regularities in our experience might be a legitimate way to establish the existence of an object, though I will not try to argue this point here.⁹

Reality and Subjective Facts

The experiential principle places strong constraints upon the truth conditions of knowable statements. (Here and in the sequel I call a statement *knowable* if it is conceptually possible that some being knows its truth value, either with certainty or not.) The principle implies that the truth value of a knowable statement can be determined from the truth values of sentences which specify what is the case for consciousness events, insofar as that truth value can be determined at all. Sentences of this latter sort (those which state what is the case for some consciousness event or other) embody or correspond to subjective facts. As in Chapter 2, let us call such sentences *subjective-fact sentences*.

A subjective-fact sentence says that something is the case for some consciousness event. By the experiential principle, the truth value of any sentence is determined, insofar as it can be determined at all, by a sufficiently long list of subjective facts. If we had enough names for consciousness events and enough sentences to express all relevant facts about how things seem, then we could express any subjective fact with a subjective-fact sentence. Then the truth value of every sentence would be determined, to the greatest extent possible, by a choice of truth-values for a sufficiently long list of subjective-fact sentences. Thus one can find a set of truth conditions which give the truth value of any sentence in terms of the truth values of subjective-fact

sentences. (By "give" I mean, of course, "determine to the extent to which it can be determined" — not "fix uniquely.")

A possible technical objection to truth conditions of this sort is their seeming circularity. The objection runs as follows. I have claimed that the truth conditions for a sentence P can be formulated in terms of the truth values of various sentences R of the form "For x it is the case that Q," where Q is some other sentence. Clearly Q contains one less phrase indicating being-the-case-for than R does. By repeating this argument, we find that the truth conditions for P depend finally upon the properties of other sentences Q which do not concern being-the-case-for. These sentences are not subjective-fact sentences, and their truth values are not fixed by those of subjective-fact sentences.

Fortunately, it is easy to rebut of this argument. One only has to recall (from Chapter 3) that the being-the-case-for operator creates nonextensional contexts. The truth value of the sentence "For x, it is the case that Q" does not depend upon the truth value of Q; hence the truth conditions of the sentence just quoted need not depend upon the actual truth value of Q. The truth conditions for "For x, it is the case that Q" can be stated in terms of *what seems to be the case* — that is, in terms of *subjective facts alone*. We already know how to evaluate such sentences in practice, using facts about how things seem to us.

In principle, one can find truth conditions for any knowable statement which make the truth value of that statement depend upon the truth values of subjective-fact statements. This is a consequence of the experiential

principle.

Idealism

Now I am ready to propose a partial answer to a traditional philosophic question about the relation between mind and reality: Does the external world which we perceive have an existence independent of our experience of it?¹⁰

We can restate this imprecise question in a more linguistic form by replacing talk about reality with talk about truth.¹¹ Then the question becomes: "Is the truth value of each statement about the real world determined completely by facts about what is experienced, or does it depend upon other facts as well?" We can rewrite the question further by replacing talk about experience with talk about subjective fact. Then the question takes this form: "Is the truth value of any statement about reality determined solely by facts about what is the case for various consciousness events?"

The preceding discussion already has answered this question in part. The truth value of any statement *whose truth value is knowable* is a function of facts about what is the case for various consciousness events. Which statements are true depends upon which facts are true for which consciousness events. For statements whose truth values cannot be known with certainty, the truth value can be fixed by such facts to the greatest extent possible. Thus, all knowable facts about reality are dependent — though in a

logical sense only — upon circumstances which are essentially mental. Any statement which might possibly be known by some subject can be true only if certain mental circumstances hold. Any statement whose truth is not a function of the mental is unknowable even in principle.

Note that this implies that a knowable statement about the objective world contains no new information that is not already in the subjective facts which determine its truth value. Once we have evaluated certain subjective-fact statements, the truth value of the other statement is fixed, at least to the extent that anyone ever can know those truth values. It is pointless to postulate an additional "something," besides the facts about how things seem, which must be real to make the statement true.

This conclusion is a version of idealism. It implies that any truth about the world which we know and experience is logically dependent upon truths about what conscious subjects are experiencing. But this idealistic thesis is far weaker than most previous versions of idealism.¹² In particular, it does not say that things depend upon mind in any manner that is causal, or that is even remotely like causality. For instance, it does not imply that things must be perceived or known in order to exist. Even less does it imply that only minds and their contents ultimately are real. All that it entails is that a knowable fact is true if and only if certain *mental circumstances* obtain. The occurrence of the mental circumstances forms a necessary and sufficient condition for the truth of the fact. In this respect, my idealism resembles the transcendental idealism of Kant.¹³

The form of idealism proposed here does not rule out the existence of knowable extramental objects, if by that one means objects which are not mental constructs or which do not depend causally upon mental activity. This idealism claims only that the facts of the knowable world are logically dependent, in a certain sense, upon the presence of consciousness in the world.

My idealism may be regarded as akin to phenomenalism, insofar as it regards the existence of physical objects as a consequence of certain facts about the possibility of experiences. However, it cannot be equated to phenomenalism of any familiar sort, since it does not attempt to reduce physical objects to anything mental — as, for example, Mill did in his equation of matter with "a Permanent Possibility of Sensation."¹⁴ In my view, matter does not consist of sensations or experiences, either of an actual or of a possible sort. Rather, the possibility of experience of a certain sort is merely a *necessary and sufficient condition* for the existence of matter. I am not suggesting that matter is "made of experience" in any sense of that phrase.

The idealism which one gets from the experiential principle does not even rule out objects which are impossible for anyone to perceive. Such an object can exist provided that the statement that it exists has truth conditions of the sort which the experiential principle requires. For example, a strictly imperceptible object might exist by virtue of the fact that the experiences of subjects always are ordered in such a way that they seem to be caused by such an object. If this is

the case, then a scientist who postulates such an object to explain certain regularities in experience will be getting the ontology right. There really would be such an object, since the regularities in experience are mental circumstances sufficient for the object's existence. The object's existence is tied to subjective facts in a way which forces the object to exist if subjects' experiences suffer from certain regularities. I am not arguing that statements about objects postulated in existing scientific theories really have truth conditions of this sort. Rather, I am introducing this possibility to show that on the idealistic view proposed here, existence cannot be reduced in any manner to being perceived. Using well-known Berkeleyan terms,¹⁵ one can say that in my idealism *esse* is very far from *percipi*.

We have arrived at this idealism by means of an argument which is essentially epistemological. This argument even has something in common with one of Berkeley's idealistic arguments — the one that begins with the observation that matter is in a sense unknowable.¹⁶ Berkeley refused to posit extramental objects, on the grounds that such objects are not needed to explain our experiences. I have no doubt that there are extramental objects, but I have refused to posit objects whose existence is not *implied by features of consciousness*. Those extra objects are not needed to explain anyone's experiences, ever. But despite the Berkeleyan flavor of my argument, my idealism is much closer to realism than Berkeley dared to go. My brand of idealism admits objective, extramental objects of perception; it is an idealism only because it postulates that the existence

of those objects depends upon the existence of consciousness in the world.

The mental circumstances which can make statements objectively true need not even be "mental" in the standard sense of that word. In Chapter 11 I argued that there may be unconscious mental processes which involve subjective fact in the same way that conscious mental processes do. The above argument for idealism does not show that conscious mental processes give rise to objective reality, as Berkeley evidently supposed sense perceptions to do.¹⁷ It shows only that subjective facts can logically "force" statements about objective reality to be true. It is conceivable that the subjective facts which do this might belong at least partly to subjects' *unconscious* psychological processes — to things that no subject knows he or she is experiencing.

All reality has the minds of individual subjects as its sources in the manner described above. However, reality does not consist solely of minds and their contents, and the minds which underpin reality need not be egos. Reality has a mental or psychological origin, but the real world is not a mere mental construct. The ultimate determinants of existence are psychological, but things can exist outside of any mind and fail to be reducible to subjective experiences. Thus, mental happenings can ensure the existence of an *extramental* world — one which cannot be said to exist only in minds.

Idealism and Physicalism

All knowable facts, including facts about the physical world, ultimately owe their truth to consciousness. The existence of physical objects and events, insofar as such phenomena are knowable, is a logical consequence of certain facts about consciousness. These are idealistic conclusions. Yet the version of idealism which I am proposing does not rule out the possibility that consciousness itself is a physical phenomenon. In fact, it is fully compatible with a *materialistic explanation of consciousness*.

Of all theories of the mind-matter relationship, idealism seems the least compatible with the view that mental processes are physical. According to idealism, matter is a product of consciousness. According to the materialist view of mind, consciousness is merely a process which takes place in a material brain. It does not seem possible for both of these views to be true.

Actually, the situation is not so simple. The particular version of idealism that I have proposed does not exclude the possibility that all mental events are events that happen to physical systems, or that consciousness has a physical explanation. Although consciousness events are the logical wellsprings of physical existence, it is logically possible that they also are events produced by the action of physical brains. I will show that this possibility does not involve any

vicious circles. An idealistic theory of reality is logically compatible with a materialistic solution to the mind-body problem. Furthermore, the possibility just mentioned has strong precedents in the literature; as I will show, ideas very much like it have been proposed both by philosophers and by noted scientists.

According to my version of idealism, the truth values of knowable statements about physical reality depend upon facts about what is the case for which consciousness event. It follows that the truth values of sentences like "John possesses a brain of such-and-such a sort in which such-and-such things are happening" also depend upon such facts. Brains owe their existence to mental circumstances, and facts about brains owe their truth to such circumstances. But this does not rule out the possibility that the consciousness events in John's history are events which happen to John's brain. Note that the circumstances which make sentences about John's brain true need not consist solely of facts about John's experiences. The experiences of other subjects would contribute as well. Consider the above quotation-marked sentence; imagine that the "such-and-suches" are filled in with real descriptions of a brain and of neural events. The mental circumstances which would make this sentence true presumably would include the fact that anyone who looks inside John's head is able to observe certain things. These circumstances also would include the performance by John of behaviors (including linguistic behaviors) which could only plausibly be explained by the presence of a brain of a certain sort. The circumstances also

might include John's having experiences of certain kinds, if these experiences can best be explained by the presence of a brain of a certain sort. The existence of John's brain, and of the events which happen to it, arise solely from mental circumstances. But none of this contradicts the possibility that certain goings-on in John's brain are themselves mental circumstances.

In this way a mental event may be physical even though the physical circumstances which explain it are products of mental circumstances. The thesis that reality arises from mental circumstances is compatible with the hypothesis that mental facts have physical explanations. Idealism, as I have formulated it here, implies only that any physical events which cause or are identical to mental events must themselves be products of mental circumstances. It has nothing to say for or against the hypothesis that mental events are simply happenings in physical systems.

The most obvious objections to the view that idealism is compatible with mind-body materialism are the threats of a causal loop and of circularity in explanation. I will address the first objection first.

The view that idealism is true but that mental events are physical seems to imply that physical events cause mental events and in turn are caused by those mental events. Actually this result does not follow, for two reasons.

The most obvious reason why there is no causal loop is that the mental events which guarantee the existence of John's brain need not be mental events belonging to John. A little earlier I listed some of the facts which could contribute

logically to the existence of John's brain and its activities. Most of these facts were not facts about what John was experiencing. (If John becomes unconscious, then his brain continues to exist. To guarantee this continuity, it is sufficient that if anyone looked inside John's head, they would find that certain tissues seem to exist there.)

A second reason why there is no causal loop is that in my idealism, mental events do not actually *cause* physical ones. The mental circumstances which underlie the being of physical events and things do not cause those events or create those things. If I observe a fire, then the fact that for my current consciousness there is a fire is at least part of the circumstances which give the fire its existence.¹⁸ But my act of observation did not set the fire; no idealist, and no arson investigator, could argue successfully that it did. The fire had its causes (say the dropping of a lighted match) which were external to my consciousness. Physical events may perhaps cause mental ones in much the same way that dropping a match causes a fire. But the act of observation which guaranteed the fire's existence did not cause the fire — and the mental events which guarantee the existence of the corresponding events in the brain need not cause those brain events.

The claim that something guarantees the existence of something else without causing it is not at all strange. A physical state of affairs can guarantee the existence of a fire without being its cause. Chemistry tells us that if electrons are being transferred from fuel atoms to oxygen atoms in a certain way and under certain conditions, then there is a fire.

But the movement of electrons described in the antecedent of the preceding sentence did not cause the fire; the dropping of the match did.

Thus one possible objection to my argument — the threat of a causal loop — is defused. I have not claimed that consciousness causes brain activity which in turn causes that same consciousness. Physical events may cause (or be) consciousness, but consciousness does not cause the physical events whose existence it underpins. The physical processes which cause (or are) consciousness events are not caused by those same consciousness events.

The second objection is that I am courting a vicious circle of explanations. Is it legitimate to explain matter as a product of mental processes, and then maintain that mental processes are goings-on in matter? Such a position looks like a viciously circular explanation — an explanation of A in terms of B and of B in terms of A, which explains nothing. Fortunately, this explanation only *seems* circular. My idealistic account of physical reality may perhaps count as an explanation of physical being (albeit an incomplete explanation), but a materialistic explanation of mind is an explanation of a fundamentally different sort. A materialistic explanation of mind ultimately has to identify mental processes with processes which happen to a physical object. Yet my idealistic explanation of matter does not identify matter with processes occurring in minds, or with anything else mental. It does not reduce physical objects and events to mental phenomena. All it claims is that the existence of physical objects and processes logically

requires, or presupposes, the occurrence of mental processes. Hence there is no vicious circularity of explanation.

It might indeed be circular to claim that physical objects and events are nothing but mental constructs, and then to claim that minds are nothing but processes occurring in material systems. ("What is matter? Just a figment of the mind. So what is mind? Just matter in a special sort of motion.") But I have not made this combination of claims. In my idealism, matter is not reducible to mind. Although there could be no matter in a consciousness-free world, material objects are not mental constructs. Hence the possibility that the mental is physical remains open.

A similar harmless apparent circularity of explanation can occur in arguments about ordinary physical events. Consider my earlier example of the fire. One might argue that a fire is reducible to certain events involving electrons in atoms and molecules. This account of fires rules out an analysis of those electronic events in terms of fires; any such account would be circular. But it does not rule out the possibility that a particular fire can guarantee the occurrence of a particular electronic event — for example, the oxidation of a particular molecule in the fuel, which event forms part of the process of burning. (This analogy is imperfect, because the fire also causes the molecular event, but it makes the point.)

My version of idealism does not preclude the physicalistic explanation of consciousness events or of experiences. It allows for the possibility that mental

processes may be events that happen to physical brains. Although the physical realm derives its being logically from the mental realm, mental events may be events in physical systems, so that a complete neurophysiological explanation of mind may exist. Well-known scientific evidence suggests that such an explanation is possible, though I think we are far from owning one. I think that the possibility of such an explanation is primarily a scientific question rather than a philosophical one.

The view that mind grounds the existence of the material world and yet still is a product of physical processes is a form of idealism quite different from any of the classical idealisms. Unlike the idealism of Berkeley or the phenomenalism of Mill, it refuses to reduce the material to the mental. Unlike the transcendental idealism of Kant, it places the self squarely in the physical world. We may call this new viewpoint *physioidealism*.

According to physioidealism, the traditional idealistic view of the physical cosmos as a product of mind is essentially correct, but it must be qualified and extended. Traditional idealistic reasoning tries to beat a path from mind to cosmos. Physioidealism suggests that such a path is too short; to fully understand the relationship between cosmos and mind, one also must take into account the origins of mind in the physical world. According to this view, the path which idealism ultimately will find is a path *from brain to cosmos*.

The concept of physioidealism has strong precedents in current philosophy and physics. A number of authors,

including both scientists and philosophers (not necessarily metaphysical idealists), have proposed ideas quite similar to it. Harold J. Morowitz has suggested that consciousness may play a fundamental role in physics but concurrently may have a natural explanation. On his view, the physical understanding of the mind may be combined with the view that mind plays an important role in quantum mechanics, to yield "an epistemological circle" which incorporates matter and consciousness.¹⁹ Erich Harth has proposed a physicalistic theory of mind, yet also has suggested that mind may play a certain role in the unification of items in the world we experience, so that mind has a crucial place in the universe.²⁰ Roger Penrose has taken a physical approach to the problems of consciousness, yet has suggested that the "actual existence" of the physical world may be related, in a way, to the existence of consciousness.²¹ John A. Wheeler has suggested that we live in "a 'participatory universe.'"²² On Wheeler's view, reality (including the past) is largely a product of observation, or of "registering" processes (which need not be conscious), even though observers and instruments are themselves parts of the universe.²³ Also, I should mention the work of various authors on the anthropic principle.²⁴ On some versions of that principle, certain properties of the physical universe are, in a sense, consequences of the presence of conscious subjects in that universe.

In my opinion, none of the earlier ideas cited in the preceding paragraph can be called explicitly and

unqualifiedly idealistic. This distinguishes them from physioidealism — though Morowitz's and Harth's ideas come quite close to physioidealism. Physioidealism acknowledges that consciousness has physical properties, but it also recognizes that the physical realm derives its being *entirely* from the mental realm. Consciousness — the occurrence of viewpoints — is the spring; the physical world is the river.

Recursive Idealism

The above discussion of the mind-matter relationship has brought us to a new view of that relationship — and of reality itself. I will now lay out this view explicitly.

The most fundamental feature of reality is consciousness. The world is first and foremost a world of conscious subjects. These subjects are genuine individuals; they are not fictions or constructs of any sort, nor are they mere composites of mental subsystems (though sometimes they behave as if they were). The fact that these subjects have the experiences they do implies the existence of an objectively real extramental world — the world of physical phenomena. That world owes its existence to the consciousness of subjects, yet it is not a mere figment or construct of minds. Physical reality is consciousness-dependent, but it is objective.

Empirical evidence drawn from neuroscience suggests that subjects are themselves products of the physical world.

The consciousness of a subject is some sort of process which happens to physical systems. Strictly speaking, empirical evidence cannot tell us whether the conscious states simply are *controlled* by happenings in the brain or actually *are* happenings in the brain. (It is easy to forget that this choice must be made on philosophical as well as scientific grounds, and to think of the mind-body problem as strictly a scientific affair.) But at very least, scientific evidence supports the view that conscious states are controlled by physical happenings in the brain.

The physical world is a logical resultant of the experiences of conscious subjects. Each physical object owes its existence to mental circumstances. Those circumstances may include, not only the fact of the object's being perceived by a subject, but also the mere *possibility* that some subject, no matter how distant at present, might perceive the object or might have its experience influenced in some way by the object.²⁵ Earlier I said that a regularity in subjects' experiences may ground the existence of an object. A tiny, almost imperceptible particle in interstellar space might never be experienced by an observer, yet it would lead to a regularity of the following sort: *if* there were an observer there with a proper detector, *then* that observer would have certain experiences. There is no compelling reason why such a counterfactual regularity cannot ground the existence of an object as surely as can an observed regularity.

A physical object, such as a bit of matter, owes its existence to experience. However, it does not owe its

existence solely to the experiences of any observer who currently is aware of it. As the above remarks about the tiny interstellar particle show, a physical object may owe its existence to its potential effects upon conscious observers who are not currently present. Thus, a physical object may ultimately owe its existence to the combined consciousnesses of *all subjects in the universe*.

A subject either is identical to a process taking place in the material world, or is linked closely to such a process. Consciousness is a product of physical phenomena. Yet the physical world is itself the byproduct of all the consciousness in it. Every conscious subject contributes, through its mental activity, to the existence of all other subjects, of itself, and of all physical items which are not subjects. A subject's present consciousness even can contribute to the existence of past and future events as well as to that of present ones. Today's memories may amount to perceptions of the past. Present historical traces and relics may make it reasonable to infer that an event happened long ago, and this may be sufficient to ensure the existence of that event (recall my remarks on imperceptible objects). That is why there can be a world before there are subjects. Past events are not brought into existence retroactively by experiences which happen now. Rather, they existed in the past because they were going to contribute to our experiences now and because they could have been experienced then if an observer had been there to do so.

Consciousness is a biological function, yet the physical world in which biological processes occur is the product of

the activity of minds. Mind grounds the being of matter, which in turn forms the medium in which minds happen. The material world gives rise to minds, and these minds are what sustain the material world in its ongoing existence.

This theory of mind-matter relations combines the chief strengths of materialism and of idealism. It allows us to maintain an idealistic view of reality — which is more satisfactory than materialism from an epistemological standpoint²⁶ — while fully acknowledging the bodily origins of mind and the successes of the scientific approach to the study of the mind. We no longer have to decide between metaphysical idealism and scientific materialism. The view of reality suggested here combines the logical virtues of idealism with the possibility of a scientific explanation of mind.

The version of idealism presented here can accurately be described as a personal idealism, since it postulates a plurality of individual subjects. It also could be labeled a process idealism, since those subjects are histories possessing genuine temporal flux. Using the terminology of the last section, we also could label it a form of physioidealism. Yet the most striking feature of this theory is its *recursive* or *self-referential* character. It uses mind to ground the existence of matter, and then declares that same matter to be the cause of mind. In the cosmos as portrayed by this theory, the experiences of each conscious subject contribute to the existence of that subject, of every other subject, and of every other object that there is. *Recursive idealism* seems a fitting name for this point of view.

The arguments for idealism given earlier in this chapter do not by themselves support recursive idealism. They lead us to a personal process idealism, but they leave open the question of whether the subject is (or is caused by) a process in physical systems. The hypothesis that subjects are physical processes, or at least are closely tied to physical processes, is supported by scientific evidence, but this hypothesis is not necessary for an idealism based on the experiential principle. Taken by itself, the experiential principle neither implies nor rules out this hypothesis. Recursive idealism is based upon a combination of the experiential principle with scientific findings.

The main lesson to be learned from recursive idealism is that scientific approaches to consciousness are compatible with the view that mind is the central feature of reality. The mere fact that every mental event has physical causes, or even that mental events are only events in utterly physical systems, cannot be used to support the view that consciousness is ontologically subordinate to matter. The discovery by scientists of a physical explanation for consciousness would not place the idealistic interpretation of reality in doubt. The possibility would remain open that the consciousness of various individual subjects ultimately gives rise to the physical conditions which cause, and even constitute, mental life.

Appendix: An Epistle from the Far Side of Relativism

As an afterthought to this discussion of recursive idealism, I will point out the implications of certain arguments in this book for philosophical relativism. The relativisms with which I am concerned are those fashionable in some quarters today, particularly on university campuses. These relativisms are doctrines which deny, in one way or another, that there is such a thing as knowable objective truth, and which attempt to replace the idea of objective truth with that of multiple truths or of multiple perspectives or viewpoints. These viewpoints may be conceived of as individual or as collective (for example, cultural, ethnic or sexual). I will not attempt to catalog all these doctrines here, since the point I will make about them is quite general.²⁷

Consider once again the argument for idealism which I completed several pages ago in the section titled "Idealism." (This was an argument for idealism as such, not for physioidealism or recursive idealism.) That argument leads to a view of reality in which every knowable object, and every knowable objective fact, is grounded entirely in subjective facts. On this view, an objective fact (or fact about the objective world) is true if and only if it follows, via some method, from some combination of subjective facts. Ultimately, a subjective fact is just a fact about how things seem. Hence an objective fact is a consequence of

facts about how things seem.

Aside from its metaphysical implications, this view of reality has profound consequences for relativism. In particular, it implies that *if relative viewpoints and perspectives exist, then there also exists an objective world.* The truth of an objective fact, if there is such a fact, depends upon the truth of a number of subjective facts. Thus, objective facts can exist or be true solely by virtue of the existence of multiple, seemingly conflicting perspectives on "reality." If one concedes that there are perspectives of this sort, or if one uses perspectives in any way in one's thinking, then one is conceding the possibility of objective truth. One cannot use perspectives in one's thinking without tacitly admitting, at very least, that the concept of objective truth makes sense.

It does not matter whether the relativism in question is a relativism of individual viewpoints, or a relativism of cultural, societal, or ethnic viewpoints. None of these viewpoints or perspectives can exist unless it is possible for things to seem some way or other. Philosophical or critical views which make use of "perspectives" always tacitly presuppose the reality of subjective fact. All discourse about perspectives is laden with subtext about subjective fact, whether those who use such discourse realize it or not.

A relativistic or anti-objective approach to knowledge implicitly acknowledges the existence, or at least the possibility, of an objective world. Within any critique of objective reality from a relativist or anti-objective standpoint, there is a subtext of subjective factuality which

undermines the overt argument of the critique. This unavoidable subtext of subjective factuality is implicitly a subtext of *objective* factuality. There is no way out of this bind in which relativism finds itself — for to abandon "subjective fact" is to abandon "perspectives."

Many versions of relativism include a critique of the notion of the conscious subject. This critique also looks different when viewed in the light of subjective fact. Consider what the world would be like if there were no objective reality, but only subjective facts — facts about how things seem in various instances. If the instances of seeming associated with those subjective facts exist for each other in the ways described in the main definitions in Chapter 5, then there are conscious subjects. Thus, there are real subjects in the world, simply by virtue of the occurrence of certain subjective viewpoints. The existence of subjects of a certain sort follows from the existence of multiple perspectives in the world, provided only that some of these perspectives are of specific kinds.

Any relativist critique of the notion of the conscious subject or self will completely miss conscious subjects of this sort. Such a critique cannot gainsay the existence of subjects without also denying that there are perspectives of a particular sort — perspectives which must exist if there is to be even the illusion of the existence of conscious subjects.

The most that a relativist critique of the subject can do is dismantle some narrower conceptions of the subject — for example, views of the subject as something introspectible, uniformly rational, or transcending the biological and social

worlds.²⁸ (I have attacked all of these views myself, especially in Chapter 10 and in the present chapter.) But no such critique of the subject can do away with the subject as such. Even a critique which purports to show that the subject is a product of social processes cannot dispose of the subject, for the possibility that the conditions for the existence of the subject are *social* cannot conflict with the fact that those conditions are *subjective*. Critiques which claim that the subject is underlain by something linguistic cannot touch the actual conscious subject — for the possibility that subjective facts are *linguistic* in character cannot alter the conditions for the existence of a conscious subject. Furthermore, the subject always is a genuine individual; a subject is not in any way conflated with other subjects, unless an actual merger or division of subjects occurs. This is the case even if the genesis of the subject is linguistic or social.

Regardless of the origins of individual consciousness, the existence of the subject, and the existence of the subject as a *true individual*, are facts safely beyond the reach of any form of relativism.

All this points to the conclusion that relativism, at least of the kind fashionable today, must be abandoned. The claim that all truth is true only from some perspective or other is so self-undermining as to be unmaintainable. Any perspective, of any kind, is a sign of the existence of an objective world. This objective world is not the property of any single individual, group, or culture; it is the joint product of all perspectives everywhere. Yet despite this diverse

From Brain to Cosmos

origin, the objective world lies outside the perspectives of any individual and of any group. Even a blatant contradiction between individual or cultural perspectives cannot compromise the reality and coherence of our common world, for somehow all perspectives manage to fit together within that shared substratum, ever partially unknown, which we call *reality*.

From Brain to Cosmos

Chapter 14

Which Beings Are Conscious?

In this chapter I will take up an important question in the philosophy of mind: Which objects in the world are conscious beings? This question is of interest for a variety of practical reasons. It is of ethical interest because of its relevance to the morality of euthanasia. It also bears on the question of the consciousness of nonhuman animals. It is of interest to the student of artificial intelligence because of its obvious connection to the question of machine consciousness.¹

Aside from practical implications, the question "Which things are conscious?" also has significance for our views about the nature of reality. Some philosophers (notably Leibniz² and more recently Charles Hartshorne³) have argued that matter is composed of conscious units. This view cannot be correct unless bits of inanimate matter can be conscious in some sense. Unless we have a way of deciding which beings are conscious and which are not, we cannot evaluate these views in any non-dogmatic manner.

In this chapter I will describe one way to determine, from empirical evidence, whether or not a given being is conscious. More precisely, I will describe a method for deciding whether a physical entity possesses a viewpoint, or a way things seem. This method is not entirely new; it can be regarded as an elaboration of certain ideas of H.S. Jennings, who was one of the pioneers of behaviorism.⁴ (However, the method neither presupposes nor leads to behaviorism.) The central idea of this method also can be found in Locke's argument for animal consciousness.⁵ The method proposed here may prove useful for deciding whether beings are conscious in some troublesome cases in which familiar criteria cannot be applied.

This goal is more modest than it might at first appear. Having a way things seem is not the only interesting feature of conscious beings. The method I will describe here does not allow us to decide whether a thing has any of the other familiar features of consciousness (such as, for example, emotion or thought). Hence it cannot decide whether a thing has consciousness of the kind which we humans would regard as normal waking consciousness. It only allows us to decide whether facts can be the case for an entity at a time — that is, whether an entity's career contains consciousness events.

Immediate Knowledge About Consciousness

Before proposing my criterion for deciding what is conscious, I will review a few of the more familiar criteria which people have used for this purpose.

In Chapter 6 I argued for the possibility of perception of mental states in other subjects. There I suggested that a subject cannot actually witness a consciousness event in another subject, but that a subject may sometimes notice that another subject is undergoing experiences of a particular sort. A consciousness event in a subject does not exist for a consciousness event in another subject, except in the extraordinary cases of merging and dividing subjects. One can know immediately about other subjects' mental states only through one's awareness of facts about those subjects — not by becoming acquainted with the other subjects' consciousness events.

Ordinarily, one knows about the mental states of other subjects after witnessing those subjects' physical behaviors. One is able to know about the inner lives of other subjects only because certain facts about those subjects are the case for one's own consciousness events. Neurologically speaking, this perception involves some sort of integration of sensations of the other subject's body, but it does not involve conscious logical inference. It seems as "immediate" as, say, the visual perception of a simple geometric pattern.

According to the ideas of Chapter 6, one can, under

certain conditions, simply notice that another being is conscious. But this method of determining what is conscious has three severe limitations.

First, this method of determining whether a being is conscious is rather limited in its applicability. It is difficult to see how to apply it to beings very different from oneself. Humans might be able to apply it to higher animals. For less humanlike beings the situation is not as clear. The kind of direct perception of mental states which I described in Chapter 6 requires a certain *rapport* between observer and observed. At very least, one must be able to recognize the behaviors of the observed being as symptoms of subjective states similar to one's own. Emotional sympathy is one case of this *rapport*; I discussed this case in detail in Chapter 6. But no suitable *rapport* seems possible between the average human and, say, the average flatworm or the average 2000-model computer. A suitable *rapport* between a human and a conscious electron (that is, an electron as Leibniz might have conceived it) seems even less likely. Even if flatworms, electrons, and simple computers were conscious, they almost certainly would lack any mental processes sufficiently humanlike to excite "gut" reactions in most humans.

A second difficulty with the method of noticing consciousness in other beings is its uselessness for ascertaining that a being is *not* conscious. One's inability to notice consciousness in another being is not sufficient evidence to conclude that that being is not conscious. This limitation of the method of noticing mental states does not

arise from the error-proneness which I discussed in Chapter 6. It is a separate difficulty.

This error-proneness which I just mentioned is the third (but most important) limitation of this method of perceiving other minds. In Chapter 6 I discussed this fallibility and identified its sources. Fallibility of this general sort occurs with all methods of knowing based on observation, but it afflicts this particular method rather severely.

Inferential Knowledge About Consciousness

Another way to establish that a being is conscious is to infer that fact logically from observed facts about that being's behavior. Philosophers have discussed this option extensively in connection with the problem of other minds.⁶ I have discussed it to some extent in Chapter 6. There are at least two known candidates for ways to decide by deliberate inference whether a being other than oneself is conscious.

One traditional solution to the problem of other minds involves inductive reasoning. I will summarize this known solution here. To deploy this solution, I begin with the knowledge that certain mental states of mine normally are associated with specific kinds of actions or behaviors. From this I infer inductively that similar actions or behaviors in other beings who resemble me also reflect mental states.⁷ This argument provides a way to find out what is conscious: simply determine which sorts of behaviors indicate consciousness in oneself (or in humans generally) and infer

that any being which exhibits such behaviors is conscious.

This known method of determining what is conscious has a well-known and crucial weakness: it depends upon inductive inferences from highly restricted classes of instances.⁸ As a solution to the problem of other minds, it asks one to generalize from a premise about oneself to a universal conclusion about all entities similar in some respects to oneself. As a way of determining whether non-human beings are conscious, it requires one to generalize further from a statement about human subjects to obtain a conclusion about all subjects. Induction from a single case or from a very special class of cases always is hazardous. But generalizing from oneself to all subjects is especially questionable because of the highly unusual character of one's experience of oneself. Confronted with the inductive solution to the problem of other minds, one is tempted to ask "How do I know that it isn't just *me* that works that way? How do I know that consciousness, as I feel and understand it, isn't just an idiosyncrasy of mine?"

Using human-like behaviors as criteria of consciousness for other beings also is suspect, in view of the very peculiar mental capacities of humans. Given the weaknesses of inferences from oneself, one obviously can criticize an inference from all humans by asking "How do I know it isn't just *humans* that work that way? How do I know that this link between behavior and consciousness is not peculiar to humans, a result of some very peculiar circuitry found only in *human* brains?"

Searle has suggested that comparisons of "the causal basis

of the behavior"⁹ of humans and animals can be used to justify the belief that a dog is conscious. As Searle recognized, this method does not share the weaknesses of a solution to the problem of other minds which relies on behavior alone.¹⁰ However (as Searle also recognized), this method does not extend easily and directly to the simpler animals.¹¹ Nor, I would add, is it useful on Leibnizian electrons. It cannot rule out *a priori* the possibility of consciousness in very simple entities, since for all we know there may be causal bases for behavior which are not like ours and yet which give rise to viewpoints.

There is another possible inferential method for deciding what is conscious. That method involves finding *logically sufficient* connections between consciousness and particular physical circumstances.¹² If one could find a physical condition which entails that an object is conscious, and if for any observed object one could determine through observation whether that condition holds, then one could decide, at least in some instances, whether an observed object is conscious. Unfortunately, this method remains far beyond our reach, given our present lack of knowledge about the nature of consciousness. Some philosophical views of mind, including materialism and epiphenomenalism, suggest that the requisite sufficient conditions might, in principle, be found. These views entail that mental processes are tied to physical ones in such a way that consciousness exists if certain physical conditions hold. If we could settle on one such theory, or could find such

conditions without presupposing any of these philosophical theories, then we might be able to construct a general test for the presence of consciousness in any being we observe. Unfortunately, we do not know of a way to do this. Hence this method of determining what is conscious remains impracticable.

Consciousness Is a Sharp Quality

All of the above methods for determining which beings are conscious appear to be either unreliable, presently impracticable, or of limited applicability. Fortunately, there is another way to determine which beings are conscious. This method relies both upon empirical facts and upon certain *a priori* principles, but it lacks many of the limitations of the other two inferential methods described above. It hinges upon one crucial property of consciousness: the fact that consciousness is, in one sense, what logicians call a *sharp property*.¹³

A sharp property is a property that is not vague — one that does not admit of degrees and borderline cases, as (for example) hotness and coldness do. If we forget about the other characteristics of consciousness and simply characterize consciousness as having a way things seem, then consciousness is a sharp property. An entity possesses consciousness now if and only if its career includes a consciousness event which occurs now — or, equivalently, if a fact is the case for that entity now. A being for which

even one fact now is the case has a "point of view" or an "inner world."¹⁴ This is true no matter how insignificant, "small," or simple the fact may be. To be genuinely unconscious, a being would have to lack *any subjective content whatsoever*. There would have to be *no* facts which are the case for it now, for even one such fact makes the being conscious. Thus, consciousness is a sharp quality in the sense that a given object either possesses it or lacks it. There are no genuine intermediate cases. Anything that has a "viewpoint," no matter how dim, or that has any subjective content, no matter how simple, is conscious without any ifs, ands, or buts, and only that which lacks all such content is nonconscious.

Dennett has argued that there is no sharp dividing line between the classes of conscious and unconscious beings.¹⁵ This claim appears to conflict with what I have just said, but actually it probably is compatible with my claim that consciousness is sharp. When I used the word "conscious," I used it in the sense of possession of subjective fact (recall Chapter 2). Dennett did not use the word in precisely this sense; hence it is possible that consciousness in my sense is a sharp property even if Dennett's claim is true. But if the two claims really are incompatible, then Dennett's claim must be rejected.

If so much as one fact is the case for an entity, then that entity has the property of being conscious. If no fact is the case for an entity, then that entity does not have the property of being conscious. Since one or the other of these two cases holds, a given entity either really possesses

consciousness or does not possess any consciousness at all. There can be no borderline cases in which something is "sort of" conscious. Stated more graphically: If one or more facts are the case for John, then John is conscious. If zero facts are the case for John, then John is not conscious. One or the other must be true.

None of this can conflict with the observed fact that consciousness has many different *levels*.¹⁶ A single subject may pass through different kinds of consciousness, such as dreaming sleep, drowsiness, and waking. (In some of these states a being appears to be only marginally conscious.) Humans and other animals may be able to have consciousness of many different kinds. One can conceive of possible beings which have consciousness of still other sorts, such as that of simple beings with "minute perceptions,"¹⁷ which come about as close to nonconsciousness as one can get without going completely blank. The experiences of such beings would have very meager content. But all of these levels and kinds of consciousness are only differences in the *quantity and quality of content*. In all of them there still are some subjective facts. As long as a being has any subjective content at all, as long as something is the case for that being at a time, then that being is conscious. Any contamination with subjective fact, no matter how marginal, removes an entity from the world of nonconscious items and places it squarely in the class of conscious beings.

This point is important enough to bear repeating. Consciousness may admit of many kinds and levels which

differ both qualitatively and quantitatively from each other. There may be vast, crucial variations in the kind and quantity of content. However, this does not imply that consciousness is *instantiated* in varying degrees. If things seem any way at all, then there is consciousness. The differences among different levels and types of consciousness are variations in the number and character of subjective facts which are the case for a subject. The *presence* of consciousness is another matter altogether. It is strictly two-valued. Either there is some subjective content for a being or there is not. If there is some such content, then the being is conscious; one can proceed to inquire as to what level or kind of consciousness it has. If there is no such content, then the being is nonconscious — like a brick in a world in which no inanimate thing is conscious.

States like somnolence, which we might colloquially describe as involving "a little bit of consciousness," are genuine cases of consciousness. Strictly speaking, they are not marginal or intermediate cases of consciousness itself. Descriptions of degrees of consciousness, like "a little bit conscious" and "fully conscious," may be useful in describing sleepy or drugged states, but they are deceptive. Such descriptions are not analogous to "a little tall" and "definitely tall." They are more closely analogous to a mathematician's descriptions of lines as "a little bit curved" or "strongly curved." All lines answering to either of these descriptions are *curved*, and that's that. They are not borderline cases between curved lines and perfectly straight Euclidean lines.

By now it should be clear what "consciousness is a sharp quality" means. A being either is conscious or is not; any state that seems to constitute a borderline case actually is a non-marginal case of consciousness which happens to involve content of an impoverished kind. A being which is conscious is radically different from one which is not. It has an "inside" (however content-rich or content-poor) as well as an "outside." A nonconscious being lacks such an "inside." For a conscious being, there truly exists (to borrow Nagel's words) "something that it is like to *be*" the being,¹⁸ no matter how rudimentary and dim that "something" is. For a nonconscious lump there is no such thing. Though admitting of degrees and levels, consciousness does not admit of borderline cases which could equally well be called cases of nonconsciousness.

The conclusion that consciousness is a sharp property actually is not as strong a thesis as it may appear. It is a consequence of the specialized way in which I have defined "consciousness." If we take "consciousness" to mean the presence of subjective fact (as I did in Chapter 2), then we find that consciousness is a sharp quality. If we had defined consciousness in a more naive way, such as the presence of sensation, feeling, and/or thought, then we would not have found this; indeed, we would have found that consciousness, like tallness, has degrees. It is important to remember that the kind of consciousness which we are investigating here is just the presence of subjective fact, as described in Chapter 2. The conclusions of this chapter will fall into better perspective if one keeps this in mind.

The Continuous Alteration Principle

The conclusion that consciousness is sharp implies a general principle which is useful for determining which beings are conscious. The following example is meant to motivate this principle. Although this example is grotesque and should not happen in reality, it is no odder from a philosophical standpoint than many personal identity examples, such as those I discussed in Chapter 12. (As I mentioned earlier, this argument, and the conclusion with which it ends, have precedents in the work of Jennings and Locke.)

Suppose that some Martians had a medical means for slightly weakening all the mental capacities of a human being. For concreteness' sake, let this be a medical procedure P which reduces one's mental and psychological powers, as quantitated on appropriate scales of measurement, by 1 percent. For example, application of P might lower intelligence just enough so that a measure of the victim's intelligence, on some scale, goes down by 1 percent; lessen emotional intensity just enough so that the response (however measured) to a given experience becomes 1 percent weaker; weaken sensation just enough so that certain measured responses to stimuli are 1 percent weaker; and so forth. Suppose that this weakening involves some actual weakening of the subjective feelings involved — not merely a measured weakening of behavior (if one thinks there is a

difference between these two kinds of weakening). I will not attempt to ask whether the subjective change is quantifiable.

Now suppose that the procedure P is performed repeatedly on a "normal" human being S. During the first application of P, all of S's psychological measurements decline by 1 percent. Clearly this change does not make S into a nonconscious lump. It merely transforms S into a slightly less intelligent and responsive person. There are many genuine conscious subjects who still rank far below S on all the scales. On the second repetition, S's already reduced mental levels decline again by 1 percent. This quantitative change also will fail to eliminate S's consciousness. S remains conscious, but the measurements now have only 0.99×0.99 , or 98.01 percent, of their initial values. Suppose, for the sake of argument, that while P is being repeatedly performed, nothing else, such as complete cessation of brain function, is allowed to happen to S. After N repetitions of P, a measurement on S whose original value was M has the value $M \times (0.99)^N$. As N becomes large, this modified value becomes smaller than the normal range of human values for the quantity being measured. However, there are many mental or psychological capacities which S never entirely loses.

If S keeps being simplified in this way, and abrupt changes (like a complete loss of mental content or death) are somehow prevented, then the final result will be something having behaviors similar in complexity to those of, say, a dog. However, the quantitative simplifications of S by

means of P must result in a being for which *some facts still are the case*. Nothing has happened that could result in the loss of all subjective facts. Let us call this new being S'. (I leave open the question of whether S' is identical in any sense to S.)

If one tried to do this experiment in practice, undoubtedly it would have a fatal outcome. However, this is not because of any conceptual impossibility in the story. A person can begin to have experiences considerably simpler and poorer than normal ones, and yet still have conscious experiences of some kind. It is likely that none of the iterations of P could change a conscious being into an entirely unconscious one. We may safely suppose that S', like S, is conscious.

An Objection Rebutted

A possible technical objection to this argument is the claim that it is a *sorites*. A sorites is a fallacious argument involving a property that seems to be preserved by the gain or loss of a small part. The classic "Tall Man" argument is a typical example of a sorites.¹⁹ This argument begins with the observation that if one reduces the height of the tall man by just a little, then one still has a tall man. By repeating this reduction, one can wrongly conclude that when only a few feet remain the man still is tall. One might think that my argument about S and S' is wrong for the same reason.

A moment's reflection shows why my argument is not a sorites. The Tall Man argument is invalid because the

property of tallness is vague; as one reduces the height of the tall man, one inevitably runs into borderline cases of tallness.²⁰ My argument about S and S' does not face the same difficulty. There are many different kinds of consciousness, and some of them can be thought of naively as degrees of consciousness. However, when it comes to the actual possession of a viewpoint, there can be no borderline cases between consciousness and utter unconsciousness. Hence consciousness lacks the vagueness upon which a sorites depends. One cannot take away consciousness entirely by reducing the amount of content by a tiny fraction. Since consciousness is sharp, the argument about P is not a sorites.

For the record, I will write out the argument about S and S' in a way which makes its validity clearer. If we let $S_0, S_1, S_2, \dots, S_N$ be the stages in the simplification of S (where S_0 is S and S_N is S'), then the argument looks like this:²¹

S_0 is conscious.

If S_0 is conscious, then S_1 is conscious.

If S_1 is conscious, then S_2 is conscious.

...

If S_{N-1} is conscious, then S_N is conscious.

Therefore, S_N is conscious.

This argument is not a sorites and is valid for any value of N . The reason that it is not a sorites is that *consciousness is sharp*.

One also might wonder whether S would ever reach a stage where the smallest possible reduction of the psychological measurements is greater than 1 percent. In that case, one can change the argument to use the minimum possible reduction. One also might wonder whether there is a stage at which any further reduction of the measured quantities would lead to zero values for some of these quantities. This circumstance might falsify an assumption which I made during the example: that subjective content is not entirely eliminated at any step. This could happen, although one might be able to postpone it by using even smaller steps, or different kinds of steps which reduce the complexity of content. Probably, the only stage at which no reduction at all is possible is the stage at which very few facts are the case for S . And a being at such a stage would be far simpler than the S' described above.

The above reply to the sorites objection underscores the meaning of the conclusion that consciousness is sharp. One can get from a tall man to a short man via a series of trifling changes; there is no distinct threshold between tallness and shortness, no sharp difference between the two. But there is a sharp logical difference between consciousness and nonconsciousness, if one defines "consciousness" as the possession of a way things seem. One can't get from this

kind of consciousness to utter non-consciousness by means of a series of insignificant mind-contractions — unless, of course, one cheats and sneaks in the removal of *all* remaining content. Such a change is not a logically necessary consequence of finite, quantitative mind-contractions. The tallness of a tall man can be pared away gradually by shrinkages of the order of millimeters. The presence of subjective fact in S cannot be pared away gradually — either S undeniably is conscious or S is not conscious. If S's mind kept on shrinking, S would remain conscious until the last trace of content was erased. At that moment, consciousness would disappear abruptly.

The Shrinking Mind Revisited

In the above example of the shrinking subject, S' exhibited behaviors far simpler than those of any ordinary human being. For argument's sake I suppose that the behavior of S' is reduced to the simplicity of a dog's behavior — that is, that all of the observable behaviors of S' are what one would expect of a dog, or are much like those of a dog in their complexity and sophistication. (For example, S' might have a poorer sense of smell and better vision than a real dog, but both S' and a typical dog exhibit sensitivity to their surroundings in qualitatively comparable degrees.)

If S is conscious, then it is plausible to suppose that S' also is conscious. If S' is conscious, and a dog has behaviors

which are similar in sophistication (if not in all details) to those of S', then it is plausible to suppose that *a dog also is conscious*.

This conclusion is not surprising. What is surprising is how we derived it.

It may be objected that the behavior of the dog is *not* a simplified version of human behavior — that one cannot get from S to the dog by making small changes, since the dog is a being of a fundamentally different kind. This objection has a point, but it still is not a successful objection. Although S' is not a dog, the differences between S' and the dog are differences that should not make the difference between consciousness and unconsciousness. These differences do not tell against the view that there is some way things seem to a particular organism. For example, the differences between a dog's senses and ours does not make or break consciousness. A normal human with full, unreduced mental capacities, who gradually lost visual acuity and acquired a keen sense of smell, would not lose consciousness by virtue of that change. A change in subjective fact would be involved, but consciousness would not be lost. One class of subjective facts would be swapped for another. The same can be said for a dog which gradually became nearly anosmic while acquiring 20-20 vision. There are other changes, besides sensory ones, which also cannot eradicate consciousness. These changes include (for example) changes in instincts and drives, in emotional responses, and in ways of interpreting the world.

One can generalize the above argument with arbitrary

pairs of beings taking the place of S and the dog. One arrives at the following general principle:

Continuous Alteration Principle (CAP). Suppose that T and U are two beings, T is conscious, and it is conceptually possible that there are beings T_1, T_2, \dots, T_N such that each being in the series T, T_1, T_2, \dots, T_N , U, besides T itself, is related to the previous member of the series by a change of one of the following two sorts:

- (1) a small quantitative change in some feature or features of behavior or of the internal processes which cause behavior;
- (2) a substitution of one kind of perception or behavior for another kind, which, if applied in a gradual manner to T or to U, would not result in T's or U's becoming unconscious.

Then it is probable that U is conscious.

As it stands, the CAP is extremely vague. It contains vague phrases, such as "small quantitative change," "kind of perception or behavior," and "gradual manner," which badly need to be assigned more precise meanings. Nevertheless, we can use the principle as a rough guide for making educated guesses about what might be conscious. And even

in its present, far-too-imprecise form, the CAP allows us to determine which beings *might* be conscious. If the CAP tells us that an entity is conscious, then we should not rule out the possibility that that being actually is conscious, since there is at least a grain of suspicion that it might be conscious.

Consciousness in the Biological World

One can use CAP as a rough-and-ready test for the presence of consciousness in a wide variety of things. In principle, one can apply CAP as a test for consciousness to animals, to elementary particles (if one has read Leibniz), to computers, or to any other observable objects one wishes to test.

The CAP strongly suggests that higher animals like dogs, horses and bats are conscious.²² This is not much of a surprise. The mental simplifications and extensions required to make humans resemble dogs (for example) seem to be conceptually possible. The CAP does not say whether the consciousness of a higher animal is much like that of a human. It does not even tell us whether higher animals have selves of a humanlike sort, since it does not tell us whether they have self-awareness.

The CAP also strongly suggests that one does not have to be a mammal to be conscious. It seems likely that it is logically possible (though not now technically possible) to alter a human being to resemble, in a general way, a frog. It

is possible to imagine intermediate states which satisfy the hypothesis of CAP. This suggests that if CAP holds then amphibians possess consciousness of some sort.

We can speculate about the possibility of consciousness in various biological systems. When we think about this possibility, we must remember that we are using "consciousness" in a very restricted sense here. To be conscious in this sense, one need not be conscious in the way that humans are conscious. One does not even have to be "awake." I am calling an object "conscious" if and only if there is a way things seem to it — that is, if and only if some fact is the case for it. Thus the claim that a frog possesses "consciousness" amounts only to this: that an instance of seeming is part of a frog's history. It does *not* imply that frogs possess intelligence of any mammal-like kind, or any thoughts, feelings, intentions, or plans such as we have. Least of all does it imply that a frog has a self, or self-awareness. All it implies is that a frog has subjective characteristics of some sort. A fact can be the case for a frog at a time.

A moment ago I said that the CAP does not imply that a frog has a self. If a frog undergoes consciousness events, then it is a conscious subject at least part of the time; hence it can be thought of as having a self in this very weak sense. But being a conscious subject is not the same as being a person (recall Chapter 5) or having a self in the usual sense of that word. If a frog were a conscious subject in the technical sense of "conscious subject" which I have used, it would not imply that a frog has a psychological self of the

sort which humans have.²³

So Which Beings Are Conscious?

Using the CAP, we have arrived at a partial, tentative answer to the question in this chapter's title. We have found reason to suppose that consciousness is a pervasive phenomenon in the world of vertebrate animals. These animals very likely possess viewpoints and undergo consciousness events of some sort.

Since the concept of consciousness which we are using is so minimalistic, it is not preposterous to ask whether invertebrates also have some kind of consciousness. I will not try to answer this difficult and provocative question here. However, I will mention some relevant facts about the behavior of simpler organisms. Invertebrates exhibit highly complex behaviors, sometimes strongly reminiscent of vertebrate behaviors.²⁴ This is true even of *unicellular* organisms, such as amoebas²⁵ and bacteria.²⁶ Bacteria, in particular, exhibit what amounts to an elementary version of memory, as well as fairly complex adaptive and goal-seeking behaviors.²⁷ Regardless of whether such organisms have experiences, it is clear that they raise interesting questions in the philosophy of mind. Philosophers would do well to study them more often.

The question of the possibility of computer consciousness is another important issue which I do not plan to take up here. However, there is no reason why the CAP could not

From Brain to Cosmos

be used to approach this problem. Since each kind of hardware and each program has its own distinctive traits, we must apply the CAP to computers on a system-by-system basis instead of trying to apply it to very wide classes of machines.

Chapter 15 Deleted

Chapter 15 has been removed from this edition of the book. In that chapter I discussed some questions in the philosophy of religion. I have done more work in that area since then, and the original chapter is no longer up to date. My later writings about the philosophy of religion are available in my other books and papers.

Chapter 16

Postscript

Now we have reached the end of the journey from experience to cosmos which we began in the first chapter of this book. We have arrived at many conclusions — most of them quite tentative — about the things, events, and persons which inhabit the world. Some of these conclusions may seem surprising, but the overall drift of our findings is even more intriguing. It is beginning to look as though some main features of the world — the kinds of things which exist, and the most crucial properties of those things — can be deduced from *facts about the nature of consciousness itself*. The existence of conscious subjects and of time are consequences of the features of conscious experience.

Further, it is beginning to look as though the external world is not as unqualifiedly external as it seems. What we have found has not cast any doubt upon the reality of the physical world; that world cannot be a mere figment of the mind, as some have speculated it might be. But the existence of the world nevertheless depends upon the

existence of consciousness in an intimate way. If the physical world contained no conscious observers at all, then there could be no physical world. Without the experiences of conscious beings, there would be no physical things to make up a cosmos. In this sense, consciousness is the source and mainspring of reality. The likelihood that consciousness is itself a phenomenon within the physical world — a natural biological process in brains — cannot change this fact.

Another consequence of our new view of reality is an updated understanding of time. We have seen that the passage of time does not involve a movement of things from the uncertainty of the future into the oblivion of the past. That movement, which plays such an important part in our everyday feelings about people and things, is merely a product of our perceptions. Nevertheless, the flow of time has a certain sort of reality. That flow is not an actual movement, but is a sort of eternal ongoing happening of all things. All events — whether past, present, and future to us now — take part in this happening. Past and future events are just as real as present ones; what is more, they are *happening*, just as are present events. We can think of the entire cosmos, encompassing all times within itself, as one vast happening, ever existent and ever in a state of flux. Although time is the fourth dimension of the physicist's world, the eternal happening of all events ensures that time is more than just an extra dimension of space.

Still another consequence of our new world-view concerns the nature of philosophy itself. We have found

what appears to be a fatal flaw in some currently popular forms of relativism. We have seen that even in a world of multiple perspectives, there can be, and indeed must be, objective realities and objective truths. Thus, the traditional philosophical quest for truth has begun to look a bit more promising than it has looked in the last few decades. The philosophical methods used in this book hint at one possible way in which this quest might be renewed and carried forward.

The ideas presented in this book do not form a closed philosophical system with a place for everything, after the manner of Hegel or Spinoza. Rather, they can be described as forming a kind of *open-ended* world system. The possibility of constructing such a system through largely analytic methods suggests that perhaps the speculative tradition in metaphysics is not as dead as many have claimed it to be, and that this tradition can be revived through a judicious choice of methods and aims.

There is much work yet to be done on the details of this project. No doubt most of the arguments I have used along the way will need to be revised or abandoned. But at least we have made a definitive start. We have found a way to explore some of the strangest questions of philosophy by beginning with what is most familiar: facts about how things seem, here and now. Perhaps more importantly, we have shown that a rational approach to the ultimate questions about our existence need not lead to ignorance or to despair. Instead, this approach can bring us to a vision of the world in which the experiences of thinking beings form as

From Brain to Cosmos

indispensable a part of the cosmic process as the hidden machinations of quarks or the vast, distant movements of the galaxies.

From Brain to Cosmos

From Brain to Cosmos

Notes

Bibliographical references, cited here by author and year, can be found in the "Works Cited" section of the book. Numbers following such citations are page numbers unless otherwise indicated.

Chapter 1. On How Things Seem to You

1. The possibility of two senses of the word "conscious" cannot be dismissed out of hand. Block has argued for a distinction between two notions of consciousness, which he calls "phenomenal consciousness and access consciousness" (Block 1996, 456). What Block calls "access consciousness" apparently does not require possession of a way things seem (Block 1996, 456-457). Searle (1992, 84), citing Block, criticizes the idea that "consciousness" has a sense of this latter kind.

2. Nagel 1974, 436.

3. Descartes, 1641.
4. Hume 1739-40, 251-253 (Book I, Part 4, Section 6); quote is on 253.
5. Russell 1912, 19. Carnap (1928, 261 (sec. 163)) also makes essentially the same point; he cites a number of precedents, including one from Russell.
6. Carnap 1928; Russell 1918, 50-51, 143-146; Russell 1924, 160-177.
7. Carnap 1928, 107-110 (secs. 67-68)).
8. Russell 1918, 143-146, especially 147.
9. See Russell 1918, 151-152.
10. An excellent introduction to phenomenology, as Husserl conceived it, is Husserl's *Cartesian Meditations* (Husserl 1950). My description of phenomenology here is based largely on Husserlian phenomenology as presented in that work. See Natanson 1973 for a general discussion of phenomenological thought, including post-Husserlian forms.
11. For Hume's terms, see for example Hume 1739-40, 1-7 (Book I, Part 1, Section 1). For Carnap's and Russell's terms, see notes 7 and 8 above, respectively.
12. I am thinking of the view that such concepts belong to "folk psychology." Searle discusses and rebuts this doctrine (Searle 1992, 5, 6, 45-48, 58-63).
13. In Dennett 1991. For example, Dennett denies the existence of qualia (Dennett 1991, Chapter 12).
14. Dennett 1991, 96-97.

Chapter 2. Into the Subjective World

1. Using Block's terminology (Block 1996, 456-457), we might call this the "phenomenal consciousness" sense.

2. Nagel 1974, 442

3. The claim which I am about to make is reminiscent of the fact that conscious experience always is associated with a particular perspective. Hence it is reminiscent of Nagel's observation that conscious experience involves a "point of view" (Nagel 1974, 437) and of Searle's emphasis on the "first-person" nature of consciousness (1992, 16, 20; see also 116-124). However, the fact which I will discuss is even more fundamental, as will become apparent later.

4. In particular, we are not asserting the existence of any mental entities of the sorts mentioned in Chapter 1.

5. On the substitutional reading of the quantifiers and its relationship to existence, see Orenstein 1978.

6. For example, D.M. Armstrong; see e.g. Armstrong 1989, 88-96.

7. Dennett 1969, pp. 8-13.

8. Husserl used "exist for" and kindred expressions in a sense not far from mine; the idea of something existing for something else occurs in Husserl's thought. See Husserl 1950, 84 (sec. 41), for one example among many. The notion of something being or existing for something can be found in Hegelian philosophy; Hegel himself used terms which have been translated as "being-for-self," "being-for-other," and "being-for-one" (Hegel 1816, 157, 119, and 159,

respectively). James, in a footnote (James 1884, 149), speaks of Hegelians who may say that "segments of the stream are consciously *for* each other." This last usage of the Hegelian terminology is closer to my conception of "existing for" or "being for."

9. For a modern approach to nonexistent objects together with discussion of older views, see Parsons 1980.

10. Nagel 1974, 436.

11. Nagel 1974, 437.

12. See Searle 1992, 16, 20; see also 116-124.

13. Ayer 1959, 67.

14. In Dennett 1991; see especially 362-367.

15. Dennett 1991, 96-97.

16. See Block 1996, especially 456-457; the kind of consciousness with which I am concerned in this book is what Block calls "phenomenal consciousness." Also see Searle 1992, 84 for a relevant (critical) remark on alternative meanings of "consciousness."

17. For a description of zombies, see for example Dennett 1991, 72-73.

18. This assessment of the theory is based on Dennett 1991; see p. 406 for the zombie-human comparison, pp. 135-137 on the role of narratives in the theory.

19. Dennett 1991, 134; see also 364.

20. Dennett 1991, 134; see also 363, 364.

21. Dennett's theory allows for borderline cases between conscious and unconscious behavior; these may give rise to the borderline cases I just mentioned. See Dennett 1991, 447.

22. Dennett seems to be suggesting the possibility of this sort of change, in Dennett 1991, 24.

23. For more information, see, for example, Dennett 1969, 20-21.

24. See Searle 1992, 159.

25. This remark is applicable to many of the various anti-philosophical or "postphilosophical" views being advocated today. For an introduction to views in this vein, see Baynes et al. 1987.

Chapter 3. The Happenings Within

1. On zombies, see for example Dennett 1991, 72-73.

2. Those familiar with the work of Carnap may detect a strong resemblance between the notion of consciousness events (especially as illustrated in this section) and Carnap's concept of "elementary experiences" (Carnap 1928, 107-110 (secs. 67-68)). Closer inspection will reveal that these two concepts are quite different. Consciousness events are *logical* entities; as I have said, they are instances of seeming, and can be thought of as features of subjective facts. Elementary experiences are not mere instances of seeming; they are products of a conceptual subdivision of "the stream of experience" (Carnap 1928, 109 (sec. 67)). Moreover, this subdivision is at least somewhat arbitrary, and is not meant to reflect any pre-existing segmentation in the stream (Carnap 1928, 109 (sec. 67)). This is not the case for consciousness events. The consciousness events in the

"stream of experience" are simply all of the instances of seeming which occur there, so there is no question of getting different consciousness events by subdividing the stream in a different way. Consciousness events are not mere segments of the stream.

3. Broad 1927, 59. (There Broad also says that psychologists use the term "*a Specious Present*" (italics Broad's)).

4. Regardless of what one thinks of its philosophical standpoint, Dennett 1991 contains an interesting discussion of several such experiments. (See especially pp. 114-115, 139-144, and 153-170; the figure of 300 milliseconds comes from p. 168, where it is given for something different from, though related to, what I am estimating).

5. Dennett (1991, 112-113) recognized that conscious processes can take significant amounts of time. (See also the preceding note.)

6. Hume 1739-40, Book I, Part IV, Sec. II (p. 207); Book I, Part IV, Sec. V (p. 233); Appendix (p. 634).

7. For a classic exposition of this view, see Boethius 524, 115-119 (Book V, Prose 6).

8. The reader unfamiliar with modal logic should consult texts in this field for the required background information.

9. Leibniz 17xx, paragraphs 53-62 (pp. 156-158), especially 56 and 62.

10. Dennett has pointed out that "a report of pain has, as it were, a built-in 'seems-to' operator" (Dennett 1969, 157). If this operator — the modality implicit in claims about how things seem — were rigorized, it would become the modal

operator which I am considering here.

11. Forbes (1985, 19) discusses "'incomplete worlds'" and "possibilities."

12. Intensionality (including that of mental sentences) is treated in, for example, Gorovitz and Williams 1969, 77-88. See also Russell 1940, 324-329.

Chapter 4. On Knowing What Just Happened

1. See Dennett 1991, 132, 317-318; see also Dennett 1969, 157, where the words "seem to seem" are employed in an example sentence. (The phrase "seem to seem" also occurs in Dennett 1991, 132.) For an illuminating fantasy involving issues of how things seem to seem, see Smullyan 1981.

2. Dennett comes close to saying this when he says "*There seems to be phenomenology.*" (Dennett 1991, 366; italics Dennett's). See also Dennett's remarks on a perceived illusory pink ring (Dennett 1991, 362-365).

3. Dennett (in Dennett 1969, 157) has used the term "'seems-to' operator" in reference to a concept which appears to be similar or identical to the one I am capturing with my "seeming operator."

4. Tense logic is discussed in (for example) Newton-Smith 1980 (especially 52-54) and Prior 1957.

5. I refer to Russell's view in Russell 1912, 19.

6. On Carruthers' view, see Carruthers 1986, 29-32. But Carruthers does give an example supporting certainty about

how things seem (1986, 30). One might wonder whether the PSR implies Carruthers' views on certainty in Carruthers 1986, although I do not think that it does.

7. With regard to this possibility and related ones, see Carruthers 1986, 28-32.

8. Russell 1912, 48 (see also 46-47).

9. For a brief description of sensory memory, see Kagan and Havemann 1976, 63-64.

Chapter 5. Conscious Beings and Their Histories

1. For background information and ideas about this problem, see for example Shoemaker and Swinburne 1989 and Hirsch 1982 (especially Ch. 10). For my understanding of this problem earlier in my career (though not for my position on it), I owe much to Shoemaker and Swinburne 1989 particularly.

2. This example is adapted from Shoemaker 1989, 86. I will discuss an example like this more thoroughly below.

3. For discussions (favorable, unfavorable, or otherwise) of such theories, see for example Carruthers 1986, 76-82; Grice 1941; Shoemaker 1970; Swinburne 1989, 8-13; Shoemaker 1989, 77-88; Hume 1739-40, Book I, Part IV, Section VI (pp. 261-262). The term "quasi-memory" is used especially in Shoemaker 1970 (272, 271 and elsewhere). Shoemaker 1989 (77-82) and Swinburne 1989 (8-11), among other authors, discuss a classic theory of this sort due

to Locke. Grice (1941, 342) discusses and rejects a view on which a kind of remembering of a state just before the present one establishes personal identity. The role of *immediately preceding experience* in this view matches that in the theory I am going to propose. The account at which Grice finally arrives in Grice 1941 is quite different from my account.

4. James 1884, 146. There are similarities between James' view of the stream of consciousness and the view I will present here. In particular, James noted that "earlier segments [of the stream] become objects for the later" (James 1884, 167, footnote). He entertained, but rejected, the view that this kind of unity of the stream simply *is* the ego (James 1884, 167, footnote); he attributed to some Hegelians a view rather similar to this view he rejected (James 1884, 149, footnote).

5. For remarks on personal identity after memory loss, see Swinburne 1989, 24-25 and Shoemaker 1989, 86-88.

6. James 1884 (though James' aim there was not to solve the problem of personal identity).

7. Foster 1979.

8. Shoemaker 1989, 86-87. For other relevant remarks on total amnesia, see Swinburne 1989, 24-25.

9. The example here is based on one from Shoemaker (1989, 87-88); I have altered some points and added the conclusion about killing. Green and Wikler (1980, 69) give a similar example, though apparently with a more thorough obliteration of brain characteristics (and with a different philosophical purpose).

10. See Swinburne 1989, 23-25, on personal identity questions about disembodiment, re-embodiment, and survival of death. On p. 25, Swinburne discusses the question of personal survival without memory.

11. *Ibid.*

12. Time as experienced by the subject of consciousness has been studied by Foster (Foster 1979, 175-176) and by Russell (Russell, 1948, 210-217), among others. Russell uses the terms "subjective time" and "objective time" (Russell 1948, 212), and refers elsewhere to "a public and a private time" (Russell 1912, 32). My ideas on the topic differ from these authors' ideas in crucial ways, though, as I have pointed out elsewhere, I owe intellectual debts to each.

13. Shoemaker 1989a, 145-147. See also Shoemaker 1989, 130-132.

14. Shoemaker 1989, 130.

15. Hirsch 1982, 286-301.

16. In Foster 1979. The quote is from p. 177.

17. The items unified into a subject history are quite different (consciousness events on my view, "presentations" on Foster's (1979, 175)), as are the relations which unify those items (continuance on my view, instead of Foster's "double overlap" (176)). My account of the subject also resembles Russell's and Carnap's views in certain respects (see chapters 1 and 3 in the present book, as well as note 32 to this chapter).

18. Relevant experiments and ideas are discussed in Dennett 1991, 114-115, 139-170.

19. See Dennett 1991, 119, 125.

20. Locke 1689, Book 2, Chap. 27 (p. 336).
21. Foster 1979, 182.
22. James 1884; the phrase itself is used on p. 146. (I should mention that James' aim in that essay was not to solve the problem of personal identity.)
23. On some psychological aspects of time, see for example Krech, Crutchfield and Livson 1969, 98, 228-229.
24. Dennett 1991, 113; see also 111-112, 253-254.
25. Dennett 1991, 111. The model is discussed in Dennett 1991.
26. Dennett 1991; particularly 135, 144, 166, 407; "presentations," 169 (see also 107).
27. Dennett 1991, 356.
28. Dennett 1991, 96-97.
29. Leibniz 17xx, paragraph 21 (p. 151). See also Leibniz 17xx, paragraphs 19-20 and 22-24 (pp. 150-151), and the modern commentary of Schrecker 1965, xv.
30. For example, Church 1956 and Drake 1974.
31. Mereology (the formal theory of wholes and parts) is discussed in an accessible way, in the context of the philosophy of mathematics, in Lewis 1991 (see especially pp. 1-3 and 72-74).
32. The logical constructions used by Russell (see for example Russell 1918, especially 143-146, and Russell 1924, 163-166) and Carnap (Carnap 1928, especially secs. 132, 136, 163) were, in my view, such substitutes. Russell's and Carnap's accounts of the self are different in central respects from mine. According to their accounts, the history of the self is a class of experiences (taken to be entities)

unified by a relation which can involve long-term memory (see Russell 1918, 148-150; Carnap 1928, pars. 78 (pp. 127-128), 108 (pp. 178-179), 120 (pp. 188-189), 132 (pp. 203-205)).

33. For a general discussion of the topology of time, covering some of the properties mentioned here, see Newton-Smith 1980, 48-54.

34. On closed time see for example Newton-Smith 1980, 57-65.

Chapter 6. Knowledge of Other Minds

1. On the link between dualism and privacy, see for example Ayer 1963, 90-91.

2. For descriptions of behaviorist views see for example Kagan and Havemann 1976, 22, and Campbell 1984, 59-64, 132-133.

3. This problem is discussed in Campbell 1984, 132-134, and (in relation to dualism) in Cornman and Lehrer 1974, 251-253, among many other works.

4. For example, in the view that dualistic interactionism raises a severe problem of other minds (see Cornman and Lehrer 1974, 251).

5. See Bloom and Lazerson 1988, 233-234.

6. See Bloom and Lazerson 1988, 234.

7. See Strawson 1959, 99.

8. Tillman 1967.

9. This is the "superactor" whom philosophers of mind

have discussed. (See Searle 1992, 35.) The actor I will discuss is not quite as convincing as the superactor.

10. Tillman 1967, especially 170-172; quote on 171.
11. Tillman 1967.
12. Tillman 1967, 171-172.
13. On the possible ways of making this step, see for example Cornman and Lehrer 1974, 251; for a particularly promising way, see Searle 1992, 71-77. On the bearing of different philosophies of mind on the problem of other minds, see for example Cornman and Lehrer 1974, 306-307, and Campbell 1984, 132-134.
14. See Ayer 1958, 249-251, and Cornman and Lehrer 1974, 251.
15. *Ibid.*

Chapter 7. The Flow of Time

1. One exemplar of the general trend of thought known as "process philosophy" is Henri Bergson, whom I will discuss shortly. The general line of thought in this chapter (aside from cited sources) owes much to Bergson 1907, although the model of time at which I finally arrive will be different. One need not accept Bergson's ideas about biology to appreciate and accept many of his ideas about time.

2. See for example McTaggart 1927, Ch. 33, pars. 333 (pp. 22-23) and 342-350 (pp. 27-31), and Mellor 1981.
3. See Bergson 1907, especially 1-9, 336-342.

4. Bergson (1907, 4) speaks of a "progress of the past which gnaws into the future."

5. Bergson 1907, 2.

6. Mellor uses the terminology of "thing-stages" in Mellor 1981, 127.

7. But see Bergson (1907, 2) for a contrasting remark on change.

8. See Bergson 1907, especially 1-9, 336-342.

9. The terminology of "object-stages" is used by Hirsch (1982; see for example 4).

10. See notes 3 and 4 above for Bergson's original version of this thought.

11. On this flow, see for example Mellor 1981, 7, 168-170, 116-118.

12. Mellor 1981, especially 7, 10, 168-170, 116-118.

13. The adjective "transitional" occurs in Whitrow 1973, (168, 175-177).

14. Richard Swinburne (Swinburne 1989, 43) makes nearly the same point when he points out that the time ordering of consecutive experiences can be "a datum of experience." This idea is implicit in Foster's view of subject identity (Foster 1979, 175-176); indeed, Swinburne (1989, 43) cites Foster as a source of his argument for this claim. Foster acknowledges that one can "directly see the movement" of a thing in motion (Foster 1979, 176).

15. See Pfeiffer et al. 1964, 163, which contains a dramatic illustration concerning frog vision.

16. Broad proposed the analogous idea that "sensible motion" implies, but is not reducible to, the succession of

apparent positions (see Broad 1927, 412). Russell (1948, 210-211) realized that motion may be noticed in one mental grasp. According to Foster's account of perception, a single "presentation" may disclose an object's being at two successive spatial positions (Foster 1979, 176).

17. On the uncertainty of time generally, see Merzbacher 1970, 25-26, 31.

18. For physical facts relevant to this phenomenon see Bueche 1986, 150-151, 224. Mellor (1981, 125) has used an example involving centripetal forces and strain in a wheel to make a point about change. However, Mellor's example is used to support a conclusion different from mine.

19. For a brief description of rolling friction, see Bueche 1986, 64.

20. For background on relativistic contraction in general, see Taylor and Wheeler 1966, 64-66; Leighton 1959, 10-11; Bueche 1986, 719-722; Einstein 1955, 34-36. On the de Broglie wavelength generally (and its smallness in the classical limit), see for example Merzbacher 1970, 2-3.

21. For definitions of instantaneous velocity and linear momentum, see e.g. Bueche 1986, 36-37, 109.

22. On these or other 4-vector quantities see Taylor and Wheeler 1966, 111-112; Leighton 1959, 28, 32-34; Einstein 1955, 44-46.

23. Bergson contrasts "a movement" with "rests placed beside rests" (1907, 312), and denies "that movement is made of immobilities" (308). See also 305-314, 336-338.

24. See note 16 above for possibly related ideas, due to Broad, Foster and Russell, about the *experience* of motion.

25. See for example Taylor and Wheeler 1966, 19.
26. Mellor 1981, 127.
27. McTaggart 1927, Ch. 33; Mellor 1981, Ch. 6.
28. For discussion and critique of this and related ideas, see Mellor 1981, especially 116, 168-170.
29. McTaggart 1927, Ch. 33; Mellor 1981, Ch. 6.
30. See Mellor 1981, 22-23 and 30, for relevant remarks on such views.
31. For a discussion of some ideas of this nature, see Newton-Smith 1980, 126-130.
32. See Mellor 1981, Ch. 6, and McTaggart 1927, Ch. 33 (especially pars. 329-333, pp. 20-23), for discussions of these logical difficulties.
33. See for example Schlesinger 1980, 36, for a passing mention of this metaphor.
34. Russell (1948, 216) proposed that "the time that occurs in the specious present is objective"; this is reminiscent of what I am proposing.
35. Whitrow 1973, 175-177.
36. See Mellor 1981, 150-155.
37. *Ibid.*
38. On connections of this sort, see Newton-Smith 1980, 11-12. For one particular idealistic perspective on these connections, see Howison 1904, especially xiii and 352.

Chapter 8. The Experience of Time

1. On this problem, see for example McTaggart 1927,

Ch. 33, pars. 329-333 (pp. 20-23), and Mellor 1981, Ch. 6. Elsewhere in this book I say more about McTaggart's and Mellor's work.

2. See for example Newton-Smith 1980, 11-12, on this issue.

3. For example, philosophers McTaggart and Mellor (cited above), who have argued that tense is inconsistent.

4. This is recognized in the theory of Mellor (1981, 78-88), for example.

5. See Mellor 1981, Ch. 6.

6. Mellor 1981, 5-6, 73-88, 29-46, 58-59.

7. Mellor 1981, 58-59, 78-88.

8. For a brief introduction to tense logic, see for example Newton-Smith 1980, 52-54, and/or Prior 1957.

9. My distinction between apparent and subjective tense is reminiscent of, though not identical to, Foster's distinction between "phenomenal and presentational time" (Foster 1979, 176; see also 177).

10. On this critique, see the sources cited in note 1 above.

11. On tense logic generally, see for example Prior 1957, especially 8-9; Forbes 1985, 38-40; Newton-Smith 1980, 52-54.

12. These two phrases, and close variants thereof, have been widely used; see, for example, Forbes 1985, 39, and Newton-Smith 1980, 53-54. Prior (1957, 9-10) explains why the phrase "It is the case that" is not similarly useful.

13. See Forbes 1985, 38-43; Newton-Smith 1980, 53-54.

14. See Forbes 1985, 38-39.

15. See Forbes 1985, 40; Newton-Smith 1980, 53-54.

16. Mellor 1981, 111.

17. Mellor (1981, 111-114) offers a rebuttal of essentially this objection. The objection, as described by Mellor, rests on the premise that colors and the like are "non-relational properties, with no more temporal than spatial connotations" (111-112). I am not concerned here with the outcome of Mellor's rebuttal as much as with the remaining intuitive difficulty of identifying properties with relations to times.

18. For an early reference on the interpretation of such terms, see Strawson 1950.

Chapter 9. Spacetime and Happening

1. Mellor 1981, 68-72.

2. Bergson (1907, 11) also noted that "The universe *endures*," and apparently held that the duration of smaller things is inherited from the universe as a whole (1907, 11). But what I have in mind is different; my idea involves *spacetime* as a whole, and requires the tenseless existence of events.

3. McTaggart 1909, 347. McTaggart used the concept of "timeless existence" (1909, 346); this may not be precisely the same idea as that of merely *tenseless* existence. I should mention also that Mellor (see Mellor 1981, for example 14-25) also uses some World War II examples, different from mine and for different purposes.

4. A number of multiple-universe concepts in current physics are discussed in Visser 1995 (4-5, 93, 249-262).

5. I am thinking, of course, of the string theories. There is a large literature, both technical and popular, on these theories. See, for example, Green, Schwarz and Witten 1987 (On higher dimensions in physics, see for example 14-16 and 25 in that reference.)

6. Broad 1927, 59-65. See also Schlesinger 1980, 31-33, 140-141.

7. Schlesinger 1980, 32.

8. Broad 1927, 64-65. Schlesinger (1980, 32) points out a possible way of blocking this regress which is different from the way I am about to present.

9. See note 2 above for a precedent in Bergson.

10. See Bergson 1907.

11. Bergson 1907, 1-10, 305-314, 336-338.

12. Minkowski 1908.

13. See Minkowski 1908, 75-80.

14. Bergson 1907, 4-5, 23. Bergson speaks of "preservation of the past in the present" (1907, 23).

15. Bergson 1907, 179, 181.

16. Bergson actually comes very close to saying this; see Bergson 1907, 11.

Chapter 10. Conscious Beings and Physical Things

1. On identity and its puzzles, see e.g. Hirsch 1982

(mostly on physical identity) and Shoemaker and Swinburne 1989 (on personal identity).

2. For a scenario of this general sort, see Hirsch 1982, 138-140.

3. Russell recognized the possibility of such an alternative description. He appeared to recognize the simplicity too, when formulating his logical atomism (Russell 1918, 143-146). Earlier, he had claimed that simplicity favors the existence of physical objects as an explanation for the regularities in experience (Russell 1912, 22-25).

4. For a closely related question (about the language used to describe the physical world), see Hirsch 1982, 138-140.

5. See Hume 1739-40, Book I, Part IV, Sec. VI (pp. 251-252) and Appendix (pp. 633-636).

6. *Ibid.*

7. See Hume 1739-40, Book I, Part IV, Sec. VI (pp. 253-254, 261-262) and Appendix (p. 635).

8. Dennett 1991, ch. 13 (especially p. 429).

9. See especially Russell 1918, 143-146, and 1924, 163-166; and Carnap 1928.

10. Mellor 1981, 127.

11. Russell's and Broad's usages of "event" seem to cover what Mellor (1981, 127) has called "thing-stages" as well as other events. See Russell 1948, 275, 305; Broad 1927, 406.

12. See Broad 1927, 393, 406-410; Russell 1948, 323, 487-488; Quine 1960, 171.

13. Mellor (1981, 17-18, 104-107) discusses this idea, and states (104) that "things are wholly present throughout

their lifetimes, and events are not."

14. See Mellor 1981, 104-105.

15. Hume 1739-40, Book I, Part IV, Sec. VI (p. 259).

16. Bergson 1907, 11, 188-189.

17. I read Hirsch 1982, for example, in this way. The views of identity proposed there do not seem to require acceptance of the view I just mentioned.

18. See Mellor 1981, 104-107, for some arguments which would refute my possibility (III) along these general lines.

19. On Hume's skepticism about the self, see Hume 1739-40, Book I, Part IV, Sec. VI (pp. 251-252); see also Appendix (pp. 633-636). For an introduction to Hume's concept of "impressions," see Hume 1739-40, Book I, Part I, Sec. I (pp. 1-7)

20. See Bergson 1907, 4-5 and 23 (on the past); 96, 179-181 (on the future); "potentialities," 179, 181.

21. Howison 1904, xiii-xiv, 352, 338-339; Bowne 1908, 143-148. (Bowne attributed "the transcendence of time" to God alone (1908, 146), but I understand this transcendence differently.)

Chapter 11. The Structure of the Self

1. For presentations and discussions of such theories, see Pears 1984 and Davidson 1982. See also Mele 1987, chs. 6 and 10, for discussion and analysis of views of this sort.

2. These problems include those revolving around the

results of split-brain surgery; on these latter problems see Marks 1980. For related or relevant ideas see Parfit 1971, and also Popper and Eccles 1985, 328-329.

3. Such "compartments" could include what are called "logic-tight compartments" (discussed in Krech, Crutchfield and Livson 1969, 766).

4. Descartes 1641, 2nd Meditation (p. 84).

5. Descartes 1641, 2nd Meditation (pp. 83-86) and 2nd and 3rd Meditations (pp. 90-91); Descartes 1637, Part Four (pp. 24-25).

6. See Descartes 1641, 6th Meditation, 133-143.

7. Plato, 439 d-e (p. 103).

8. Hartmann 1868, for example part I, Introductory, Sec. I (pp. 3-5); Schopenhauer 1844, especially vol. II, Chap. 19.

9. See Pears 1984 (especially Ch. 5) and Davidson 1982.

10. See for example Marks 1980. For related or relevant information and ideas see Popper and Eccles 1985, 311-333 (especially 329).

11. See for example Dennett 1991, 14, 259-263.

12. See, for example, Kagan and Havemann 1976, 379.

13. Marks (1980, 17) considers split-brain cases in which "simultaneous conscious experiences" are not noticed at once in a single mental act.

14. Block has suggested (Block 1996, 457) that the contents of the Freudian unconscious might be instances of what he calls "phenomenal consciousness" (1996, 456). This suggestion amounts to the same thing as I am proposing here.

15. For definitions of the relevant terms see Goldenson

1984, 173 and 771. The definitions I have used may not be exactly equivalent to these.

16. See Goldenson 1984, 597, for a definition. (This definition may not be exactly equivalent to mine.)

17. See note 1 above for a reference on such views.

18. Pears 1984, 67.

19. See Quine 1959, 231-232, for the method I have in mind.

20. The concept of "logic-tight compartments" is discussed briefly in Krech, Crutchfield and Livson 1969, 766.

21. Carnap anticipated this view in a way; he held that the self is both a "unit" and a "class of elementary experiences" (1928, 260 (sec. 163); italics removed from second quote). However, Carnap viewed the self as something abstract — specifically, a class, which for Carnap is a unity of sorts.

Chapter 12. Personal Identity: Some Problems

1. A good introduction to questions like this is Shoemaker and Swinburne 1989. My way of presenting these problems owes much to that work, but its authors should not be blamed for my conclusions.

2. On split-brain operations, see for example Marks 1980 (especially 1-6); Popper and Eccles 1985, 311-329. My general line of interest in these operations owes much to Marks' book, which discusses and analyzes the idea "that the

split-brain patient has two minds" (1980, 1).

3. On dividing and/or fusing persons, see for example Parfit 1971, 4-7; Shoemaker 1970, 278-280; Swinburne 1989, 14-16, 21, 45; Shoemaker 1989, 84-85.

4. See note 2 above for some references on split-brain surgery.

5. See Parfit 1971; Swinburne 1989, 14-16; Shoemaker 1989, 84-85. See also Shoemaker 1970, 278.

6. Foster 1979, 182.

7. See James 1884 and Foster 1979 (especially p. 176) for viewpoints which cohere with this intuition.

8. Black 1952, 253-262.

9. This principle is described in (for example) Loux 1970, 236; see also Black 1952.

10. Marks (1980) has mentioned an idea which can be considered a variant of this — namely, that a person (ordinarily so called) might be, in some sense, two conscious beings. (See Marks 1980, 7, 35.)

11. Marks (1980) considers problems of the identity of the mind in split-brain cases, and considers the question of which splitting product is the original person's mind (p. 9).

12. For this or related background information, see for example Marks 1980, 1, 5, 8; Popper and Eccles 1985, 350-354.

13. Popper and Eccles 1985, 350-354.

14. Popper and Eccles 1985, 325-326, 315, 331. See also 328-330.

15. Popper and Eccles 1985, 315, 330-333; see also 316-329. (However, Eccles did not advocate the view that the

minor hemisphere has its own consciousness (Popper and Eccles 1985, 328).)

16. See Marks 1980, 17-19, 26-28, for relevant facts about these influences.

17. The view that split-brain patients are in some sense double persons has been debated by philosophers. See Marks 1980 on this and related issues.

18. See Popper and Eccles 1985, 315-329 (especially 328).

19. See for example Swinburne 1989, 21, 45.

20. A more specific scenario for splitting-and-fusing beings is found in Parfit 1971, 22-23.

21. See e.g. Mellor 1981, 104-107, for a view like this regarding people (and a mention of the prephilosophic view).

Chapter 13. Mind and Matter

1. The major kinds of theories of truth are discussed in Grayling 1982, Chapters 5 and 6, which contains background information used in this chapter.

2. This ambiguity was noted by Tarski (1944, 342).

3. Hartshorne, a believer in God, has made what amounts to the suggestion that the existence of God cannot, strictly speaking, be regarded as a fact (1962, 296). He also has written that "for religion, God is a principle and not a mere fact" (Hartshorne 1965, 126). However, what my Bob character has in mind is more prephilosophical in character.

The ideas of my Bill and Bob characters are not based on Hartshorne's theories.

4. For a reference on theories of truth in general, see note 1 above.

5. On the semantic definition or theory of truth see Tarski 1944 and Tarski 1931; also Grayling 1982, 157-163.

6. See Drake 1974, 97, for the real result which I am paraphrasing here.

7. An example is the pure first-order quantifier calculus; see Church 1956, 246. (Church 1956 explains the concepts involved far more rigorously than I have done, and gives the exact result which I am loosely paraphrasing here.)

8. The arguments in this section owe something to Berkeley's well-known arguments for idealism (see Berkeley 1710 and Berkeley 1713). However, my conclusions will be quite un-Berkeleyan in several respects.

9. Kant seems to have allowed for a possibility much like this. See Kant 1781, "Transcendental Logic," 1st Division, Book II, Ch. II, Sec. 3, pt. 4 (especially pp. 190-193). Mill's view of matter as "a Permanent Possibility of Sensation" (Mill, 243) certainly implies this or something close.

10. The argument and conclusion which I am about to present have precedents in Kant, Mill and Berkeley. I will discuss these precedents later.

11. For ideas distantly relevant to this replacement see for example Quine 1939.

12. Traditional idealisms and weaker versions are contrasted in Grayling 1982, 280-288, where some idealistic concepts quite different from mine are discussed.

13. See Kant 1781, especially "Transcendental Logic," 1st Division, Book II, Ch. II, Sec. 3, pt. 4 (pp. 190-193).
14. Mill 1865, 243.
15. Berkeley used the two italicized terms in Berkeley 1710, part I, par. 3 (p. 23).
16. Berkeley 1710, Part I, pars. 18-20 (pp. 29-31).
17. See especially Berkeley 1710, Part I, par. 1 (p. 22).
18. I should mention that I thought of this example after reading a very different example of Mellor's (1981, 177-178), which is about an entirely different topic (not related to idealism) and has a different conclusion unrelated to mine. The two examples share only the mention of fire, matches, and causal loops.
19. Morowitz 1980, 39-40.
20. Harth 1993, 7-10 and 172-173. (Harth's theory of mind is discussed in Harth 1993.)
21. Penrose 1989, 448. (The physical approach to consciousness is found in Penrose 1989.)
22. Wheeler 1983, 194.
23. Wheeler 1983, 194-199; the word "registering" is found on 194.
24. Brief discussions of the anthropic principles as they relate to problems of consciousness are found in Penrose 1989, 433-434, and in Harth 1993, 12-14.
25. See the sources cited in note 9 above for precedents to this view in Kant and Mill.
26. Berkeley already recognized this advantage of idealism; see for example Berkeley 1710, part 1, par. 88 (65).

27. This summary of the current wave of relativism is based on a wide range of literature and discussions. More sophisticated ideas on the so-called "end of philosophy" are discussed in the General Introduction to Baynes et al. 1987 and/or represented in some of the articles in that work.

28. See the General Introduction to Baynes et al. 1987 (especially p. 4) for a discussion of recent criticisms of these ideas of the subject.

Chapter 14. Which Beings are Conscious?

1. There is a large current literature on the prospects for machine consciousness and other mental features of machines. Interesting older sources include these articles in Hook 1960: Danto 1959; Hook 1959; Lachman 1959; Scriven 1959; Watanabe 1959; Weiss 1959.

2. Leibniz 17xx, paragraphs 66-70 (p. 159).

3. See Hartshorne 1984, 60-63 and Hartshorne 1962, 191-196.

4. A very interesting older discussion of invertebrate behavior is found in Jennings 1906. Jennings, a pioneer of behaviorism (see Jensen 1962, xvi), noticed the "continuity" (Jennings 1906, 335) of behavior through the biological world, and thought it consistent with the hypothesis of invertebrate consciousness (336). The continuous alteration principle, which I will introduce below, can be regarded as a sharpened version of his ideas about this continuity. For some comments on Jennings' work, see Jensen 1962.

5. Locke 1689, Book 2, Chapter 9 (pp. 148-149).

6. See for example Ayer 1958, 243-254. My discussion in the next two paragraphs owes much to these comments of Ayer's (particularly Ayer 1958, 249-250). I discussed the problem of other minds in Chapter 6.

7. Ayer 1958, 249. (My version of this argument, in this sentence and the last, follows Ayer's version closely, though perhaps not perfectly.)

8. Ayer 1958, 249-250.

9. Searle 1992, 73.

10. Searle 1992, 71-74, 21-22.

11. Searle 1992, 74.

12. See Searle 1992, 74-75, for a similar idea not involving logical entailment.

13. Searle (1992, 83) has noted what appears to be the same fact; he has pointed out that "a system is either conscious or not" even though it can have "different degrees of consciousness."

14. Hume's and Locke's oyster examples come to mind at this point (Hume 1739-40, Appendix, p. 634; Locke 1689, Book 2, Chapter 9, p. 148). The expression "point of view" is used, in a sense rather similar to the one I have in mind, by Nagel (1974, 437).

15. Dennett 1991, 447.

16. Searle also recognized that this fact is reconcilable with the fact that "a system is either conscious or not" (Searle 1992, 83).

17. Leibniz 17xx, par. 21 (p. 151); see also pars. 19-20 and 22-24 (pp. 150-151), and Schrecker 1965, xv.

18. The quoted phrase is from Nagel's characterization of conscious organisms in Nagel 1974, 436.

19. See Forbes 1985, 164, 171-172. I have drawn on Forbes 1985 for information about sorites arguments in general.

20. See Forbes 1985, 164, 169.

21. The presentation of the Tall Man argument in Forbes 1985 (172) has the same form, though of course the subject matter is different.

22. Nagel's suggestion (in Nagel 1974, especially pp. 439-440) that a bat's experience is in some respects unimaginable to us does not threaten this conclusion about bats!

23. Dennett has argued that a bat lacks a significant "selfy self" (Dennett 1991, 448).

24. See note 4 above for references on the work of Jennings, who explored invertebrate behavior, noted the similarities to vertebrate behavior, and anticipated some of the ideas I will present in this chapter.

25. See Jennings 1906, 1-25.

26. See Alberts et al. 1983, 757-763, on a bacterial behavior which can be viewed in this way. See also Jennings 1906, 26-40.

27. I refer to *chemotaxis*, which is discussed in Alberts et al. 1983 (757-763), and which is referred to in the following title of a section of a chapter: "Bacterial Chemotaxis Is a Simple Kind of Intelligent Behavior" (p. 757).

Note on References and Works Cited

The references for Chapter 15 were removed from the book when Chapter 15 was removed.

The Works Cited list is the same as in the original edition of the book. Some of the items in the Works Cited list were cited only in Chapter 15. I left these items in the list.

Works Cited

This list contains all works used as sources of information or ideas in this book. It is not a comprehensive bibliography of any sort. Many of the topics discussed in this book are subjects of vast bodies of published literature; others, such as introductory physics, are covered in many good books. In cases of these sorts, I concentrated on typical reference sources which I felt would be useful to the reader, or which I personally found helpful. (In areas of active research, these may not be the most current works available.) No slight is intended toward any work not mentioned in this list.

Dates following author's names are meant to be (approximate) publication dates unless a separate publication date is given, in which case they are meant to be (approximate) dates of first publication or creation. The latter dates come from the works themselves or their front matter, or occasionally from Durant 1953. Dates listed in this section should not be treated as exact; some may be educated guesses.

From Brain to Cosmos

- Alberts, B., D. Bray, J. Lewis, M. Raff, K. Roberts, and J. D. Watson 1983. *Molecular Biology of the Cell*. New York: Garland Publishing.
- Armstrong, D.M. 1989. *Universals: an Opinionated Introduction*. Boulder, Colorado: Westview Press.
- Ayer, A. J. 1958. *The Problem of Knowledge*. London: Macmillan & Co.
- 1959. Privacy. Reprinted in *The Concept of a Person and Other Essays*. N.Y.: St. Martin's Press, 1963.
- 1963. The concept of a person. *The Concept of a Person and Other Essays*. N.Y.: St. Martin's Press.
- Baynes, K., Bohman, J., and McCarthy T. (eds.) 1987. *After Philosophy: End or Transformation?* Cambridge, Mass.: The MIT Press.
- Bergson, H. 1907. *Creative Evolution*. Trans. A. Mitchell. N.Y.: Henry Holt and Co., 1923.
- Berkeley, G. 1710. *A Treatise Concerning the Principles of Human Knowledge*. In *Principles, Dialogues, and Philosophical Correspondence*, ed. and intro. by C. M. Turbayne. Indianapolis: Bobbs-Merrill Educational Publishing, 1965.
- 1713. *Three Dialogues Between Hylas and Philonous*. In *Principles, Dialogues, and Philosophical Correspondence*, ed. and intro. by C. M. Turbayne. Indianapolis: Bobbs-Merrill Educational Publishing, 1965.

- Black, M. 1952. The identity of indiscernibles. In *Universals and Particulars: Readings in Ontology*, ed. M. J. Loux. Revised ed. Notre Dame: Univ. of Notre Dame Press, 1976.
- Block, N. 1996. How can we find the neural correlate of consciousness? *Trends in Neurosciences* (Reference Edition) 19:456-459.
- Bloom, F. E. and A. Lazerson 1988. *Brain, Mind, and Behavior*. 2nd ed. New York: W. H. Freeman & Co.
- Boethius, 524. *The Consolation of Philosophy*. Trans., with intro. and notes, by R.H. Green. Indianapolis: The Bobbs-Merrill Co., 1962.
- Bowne, B. P. 1908. *Personalism*. Boston: Houghton Mifflin Co.
- Broad, C. D. 1927. *Scientific Thought*. Reprint ed. New York: Harcourt, Brace & Co.
- Bueche, F. J. 1986. *Introduction to Physics for Scientists and Engineers*. 4th ed. N.Y.: McGraw-Hill.
- Campbell, K. 1984. *Body and Mind*. 2nd ed. Notre Dame, Ind.: Univ. of Notre Dame Press.
- Carnap, R. 1928. *The Logical Structure of the World*. In *The Logical Structure of the World; Pseudoproblems in Philosophy*, trans. R. A. George. Berkeley: University of California Press, 1969.
- Carruthers, P. 1986. *Introducing Persons*. London: Croom Helm.
- Church, A. 1956. *Introduction to Mathematical Logic*. Vol. 1. Princeton, N.J.: Princeton Univ. Press.
- Cornman, J. W. and K. Lehrer, 1974. *Philosophical*

From Brain to Cosmos

- Problems and Arguments: An Introduction.* 2nd ed. N.Y.: Macmillan Pub. Co.
- Danto, A. C. 1959. On consciousness in machines. In Hook 1960.
- Davidson, D. 1982. Paradoxes of irrationality. In *Philosophical Essays on Freud*, ed. R. Wollheim and J. Hopkins. Reprint ed. Cambridge, England: Cambridge Univ. Press, 1982.
- Dennett, D. C. 1969. *Content and Consciousness*. London: Routledge & Kegan Paul.
- 1991. *Consciousness Explained*. Boston: Little, Brown and Co.
- Descartes, R. 1637. *Discourse on the Method of Rightly Conducting the Reason and Seeking Truth in the Sciences*. In *Discourse on Method and Meditations*, trans. and intro. by L.J. Lafleur. Indianapolis: The Bobbs-Merrill Co., 1960.
- 1641. *The Meditations Concerning First Philosophy*. In *Discourse on Method and Meditations*, trans. and intro. by L.J. Lafleur. Indianapolis: The Bobbs-Merrill Co., 1960.
- Drake, F. R. 1974. *Set Theory: An Introduction to Large Cardinals*. Amsterdam: North-Holland Pub. Co.
- Dummett, M. 1964. Bringing about the past. In *Truth and Other Enigmas*. Cambridge, Mass.: Harvard Univ. Press, 1978.
- Durant, W. 1953. *The Story of Philosophy*. 2nd ed. N.Y.: Simon and Schuster.
- Einstein, A. 1955. *The Meaning of Relativity*. 5th ed.

From Brain to Cosmos

- Trans. E. P. Adams, E. G. Straus, and S. Bargmann.
Princeton, N.J.: Princeton Univ. Press.
- Findlay, J. N. 1948. Can God's existence be disproved?
Mind 57:176-183.
- Forbes, G. 1985. *The Metaphysics of Modality*. Oxford:
Clarendon Press.
- Foster, J. 1979. In *self-defence*. In *Perception and Identity*,
ed. G. F. Macdonald. London: The Macmillan Press Ltd.
- Goldenson, R. M., ed. (Editor-in-Chief), 1984. *Longman
Dictionary of Psychology and Psychiatry*. N.Y.:
Longman.
- Gorovitz, S., and R. G. Williams, in collaboration with D.
Provence and M. Provence, 1969. *Philosophical Analysis:
An Introduction to Its Language and Techniques*. 2nd ed.
N.Y.: Random House.
- Grayling, A. C. 1982. *An Introduction to Philosophical
Logic*. Sussex: The Harvester Press Limited; Totowa,
N.J.: Barnes & Noble Books.
- Green, M. B., J. H. Schwarz, and E. Witten, 1987.
Superstring Theory. Vol. 1. Cambridge, England:
Cambridge Univ. Press.
- Green, M. B., and D. Wikler, 1980, Brain death and personal
identity. Reprinted in *Medicine and Moral Philosophy*,
ed. M. Cohen, T. Nagel and T. Scanlon. Paperback ed.
Princeton, N.J.: Princeton Univ. Press, 1982.
- Grice, H. P. 1941. Personal identity. *Mind* 50:330-350.

- Harth, E. 1993. *The Creative Loop: How the Brain Makes a Mind*. Reading, Mass.: Addison-Wesley Pub. Co.
- Hartmann, E. v. 1868. *Philosophy of the Unconscious*. Trans. W. C. Coupland. N.Y.: Harcourt, Brace & Co. 1931.
- Hartshorne, C. 1962. *The Logic of Perfection and Other Essays in Neoclassical Metaphysics*. La Salle, Ill.: Open Court.
- 1965. *Anselm's Discovery: A Re-examination of the Ontological Proof for God's Existence*. La Salle, Ill.: Open Court.
- 1984. *Omnipotence and Other Theological Mistakes*. Albany, N.Y.: State Univ. of New York.
- Hegel, G. W. F. 1816. *Hegel's Science of Logic*. Trans. A. V. Miller. Reprint ed. N.Y.: Humanities Press, 1976.
- Hirsch, E. 1982. *The Concept of Identity*. N.Y. and Oxford: Oxford University Press.
- Hook, S., ed., 1960. *Dimensions of Mind*. N.Y.: New York Univ. Press.
- Hook, S. 1959. A pragmatic note. In Hook 1960.
- Howison, G. H. 1904. *The Limits of Evolution and Other Essays*. 2nd ed. (rev.) N.Y.: The Macmillan Co.
- Hume, D. 1739-40. *A Treatise of Human Nature*. Ed. and indexed by L. A. Selby-Bigge; revised by P. H. Nidditch. Reprint of 2nd ed. Oxford: Clarendon Press, 1983.
- Husserl, E. 1950. *Cartesian Meditations: An Introduction to Phenomenology*. Trans. Dorion Cairns. Dordrecht: Kluwer Academic Publishers, 1993.
- James, W. 1884. On some omissions of introspective

- psychology. In *Essays in Psychology*. Textual editor, Fredson Bowers. Cambridge, Mass.: Harvard Univ. Press, 1983.
- Jennings, H. S. 1906. *Behavior of the Lower Organisms*. Intro. D. D. Jensen. Bloomington: Indiana Univ. Press, 1962.
- Jensen, D. D. 1962. Foreword to the 1962 edition. In Jennings 1906.
- Kagan, J., and Havemann, E. 1976. *Psychology: An Introduction*. 3rd ed. N.Y.: Harcourt Brace Jovanovich.
- Kant, I. 1781. *Critique of Pure Reason*. Revised and expanded from translation by J. M. D. Meiklejohn. Ed. Vasilis Politis. Everyman Library (London: J. M. Dent; Vermont: Charles E. Tuttle), 1993.
- Kenny, A. 1979. *The God of the Philosophers*. Oxford: Clarendon Press.
- Krech, D., Crutchfield, R. S., and Livson, N. 1969. *Elements of Psychology*. 2nd ed. N.Y.: Alfred A. Knopf.
- Lachman, R. 1959. Machines, brains, and models. In Hook 1960.
- Leibniz, G.W.v. 17xx. *Monadology*. In *Monadology and Other Philosophical Essays*, trans. P. Schrecker and A. M. Schrecker; intro. and notes by P. Schrecker. Indianapolis: The Bobbs-Merrill Company, 1965.
- Leighton, R. B. 1959. *Principles of Modern Physics*. N.Y.: McGraw-Hill Book Co.
- Lewis, D. 1991. *Parts of Classes*. Oxford: Basil Blackwell.
- Locke, J. 1689. *An Essay Concerning Human*

- Understanding*. Ed. and foreword by P.H. Nidditch.
Reprint paperback ed. Oxford: Clarendon Press, 1979.
- Loux, M. 1970. Particulars and their individuation. In
Universals and Particulars: Readings in Ontology, ed. M.
J. Loux. Revised ed. Notre Dame: Univ. of Notre Dame
Press, 1976.
- Malcolm, N. 1960. Anselm's ontological arguments. *The
Philosophical Review* 69:41-62.
- Marks, C. E. 1980. *Commissurotomy, Consciousness, and
Unity of Mind*. Montgomery, Vt.: Bradford Books.
- McTaggart, J. E. 1909. The relation of time and eternity.
Mind 18:343-362.
- 1927. *The Nature of Existence*. Ed. C. D. Broad. Vol.
2. Cambridge: Cambridge Univ. Press, 1927
- Mele, A. R. 1987. *Irrationality: An Essay on Akrasia, Self-
Deception, and Self-Control*. N.Y.: Oxford University
Press.
- Mellor, D.H. 1981. *Real Time*. Cambridge, England:
Cambridge Univ. Press.
- Merzbacher, E. 1970. *Quantum Mechanics*. 2nd ed. N.Y.:
John Wiley & Sons.
- Mill, J.S. 1865. *An Examination of Sir William Hamilton's
Philosophy, and of the Principal Philosophical Questions
Discussed in His Writings*. Vol. I. Boston: William V.
Spencer.

From Brain to Cosmos

- Minkowski, H. 1908. Space and time. In *The Principle of Relativity*, trans. W. Perrett and G. B. Jeffery. N.Y.: Dover Publications, 1952.
- Morowitz, H.J. 1980. Rediscovering the mind. Reprinted in *The Mind's I*, composed and arranged by D.R. Hofstadter and D.C. Dennett. N.Y.: Basic Books, 1981.
- Nagel, T. 1974. What is it like to be a bat? *The Philosophical Review* 83:435-450.
- Natanson, M. 1973. *Edmund Husserl: Philosopher of Infinite Tasks*. Evanston: Northwestern University Press.
- Newton-Smith, W.H. 1980. *The Structure of Time*. Paperback ed. London: Routledge & Kegan Paul, 1984.
- Orenstein, A. 1978. *Existence and the Particular Quantifier*. Philadelphia: Temple Univ. Press.
- Parfit, D. 1971. Personal identity. *The Philosophical Review* 80:3-27.
- Parsons, T. 1980. *Nonexistent Objects*. New Haven: Yale Univ. Press.
- Pears, D. 1984. *Motivated Irrationality*. Oxford: Clarendon Press.
- Penrose, R. 1989. *The Emperor's New Mind*. N.Y.: Oxford University Press.
- Pfeiffer, J., and the editors of *Life*, 1964. *The Cell*. N.Y.: Time Inc.
- Plantinga, A. C. 1977. *God, Freedom, and Evil*. Reprint ed. Grand Rapids, Mich.: William B. Eerdmans Pub. Co., 1989.

- Plato. *Plato's Republic*. Trans. G.M.A. Grube;
Indianapolis: Hackett Pub. Co., 1974.
- Popper, K. R. and Eccles, J. C. 1985. *The Self and Its Brain*.
2nd printing (corrected). Berlin: Springer International.
- Prior, A. N. 1957. *Time and Modality*. Oxford: Clarendon
Press.
- Quine, W. V. O. 1939. A logistical approach to the
ontological problem. In *The Ways of Paradox and Other
Essays*. Revised & enlarged ed. Cambridge, Mass.:
Harvard Univ. Press, 1976.
- 1959. *Methods of Logic*. Rev. ed. N.Y.: Henry Holt
& Co.
- 1960. *Word and Object*. Cambridge, Mass.: The
Technology Press of the Massachusetts Institute of
Technology; New York and London: John Wiley & Sons.
- Rescher, N. 1959. The ontological proof revisited.
Australasian Journal of Philosophy 37:138-148.
- Russell, B. 1912. *The Problems of Philosophy*. Reprint
paperback ed. London: Oxford Univ. Press, 1976.
- 1918. "The Philosophy of Logical Atomism." In *The
Philosophy of Logical Atomism*, ed. and intro. by D. Pears.
La Salle, Ill.: Open Court, 1985.
- 1924. "Logical Atomism." In *The Philosophy of
Logical Atomism*, ed. and intro. by D. Pears. La Salle, Ill.:
Open Court, 1985.
- 1940. *An Inquiry into Meaning and Truth*. N.Y.: W.
W. Norton & Co.
- 1948. *Human Knowledge: Its Scope and Limits*.
N.Y.: Simon and Schuster.

- Schlesinger, G. N. 1980. *Aspects of Time*. Indianapolis: Hackett Pub. Co.
- Schopenhauer, A. 1844. *The World as Will and Representation*. Trans. E.F.J. Payne. Vol. 2. N.Y.: Dover Publications, 1966.
- Schrecker, P. 1965. The unity of Leibniz' philosophic thought. In Leibniz, G. W., *Monadology and Other Philosophical Essays*, trans. P. Schrecker and A. M. Schrecker; intro. and notes by P. Schrecker. Indianapolis: The Bobbs-Merrill Company, 1965.
- Scriven, M. 1959. The compleat robot: a prolegomena to androidology. In Hook 1960.
- Searle, J.R. 1992. *The Rediscovery of the Mind*. Cambridge, Mass.: The MIT Press.
- Shoemaker, S. 1970. Persons and their pasts. *American Philosophical Quarterly* 7:269-285.
- 1989. Personal identity: a materialist's account. In Shoemaker and Swinburne 1989.
- 1989a. Sydney Shoemaker's reply. In Shoemaker and Swinburne 1989.
- Shoemaker, S. and Swinburne, R. 1989. *Personal Identity*. Reprinted (first published 1984). Oxford: Basil Blackwell.
- Smullyan, R.M. 1981. An epistemological nightmare. In *The Mind's I*, composed and arranged by D.R. Hofstadter and D.C. Dennett. N.Y.: Basic Books, 1981.
- Spinoza, B. 1677. *Ethic*. Trans. W.H. White. In *Spinoza Selections*, ed. John Wild. N.Y.: Charles Scribner's Sons, 1930.

- Strawson, P. F. 1950. On referring. *Mind* 59:320-344.
- 1959. *Individuals: An Essay in Descriptive Metaphysics*. London: Methuen & Co.
- Stump, E. and Kretzmann, N. 1981. Eternity. *The Journal of Philosophy* 78:429-458.
- Swinburne, R. 1989. Personal identity: the dualist theory. In Shoemaker and Swinburne 1989.
- Tarski, A. 1931. The concept of truth in formalized languages. In *Logic, Semantics, Metamathematics*, trans. J.H. Woodger. Oxford: Clarendon Press, 1956.
- 1944. The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research* 4:341-376.
- Taylor, E. F. and J.A. Wheeler 1966. *Spacetime Physics*. San Francisco: W. H. Freeman and Co.
- Tillman, F. A. 1967. On perceiving persons. In *Phenomenology in America*, ed. and intro. by J. M. Edie. Chicago: Quadrangle Books, 1967.
- Visser, M. 1995. *Lorentzian Wormholes: From Einstein to Hawking*. Woodbury, NY: American Institute of Physics.
- Watanabe, S. 1959. Comments on key issues. In Hook 1960.
- Weiss, P. 1959. Love in a machine age. In Hook 1960.

From Brain to Cosmos

- Wheeler, J.A. 1983. Law without law. In *Quantum Theory and Measurement*, ed. J. A. Wheeler and W. H. Zurek. Princeton, N.J.: Princeton Univ. Press, 1983.
- Whitrow, G. J. 1973. *The Nature of Time*. N.Y.: Holt, Rinehart and Winston.

Appendix A

Notes on *From Brain to Cosmos*:
Questions and Answers about Subjective Fact

Mark F. Sharlow

In my book *From Brain to Cosmos*, I made extensive use of the concepts of a subjective fact and of a subjective fact statement. Here I will try to answer some questions and objections regarding these concepts. These questions did not come from readers of the book, but are questions that I anticipate some readers might have.

Question 1. Are subjective facts theory-laden?

Question 2. If subjective facts were theory-laden, would this undermine the project of *From Brain to Cosmos*?

Question 3. Is the language of subjective fact an independent observation language?

Brief Answers to Questions 1-3:

1. Yes, in a certain sense.
2. No.
3. No.

Longer Answers to Questions 1-3:

Philosophers of science have long recognized that observation statements can be "theory-laden," or dependent in a certain way upon theoretical presuppositions. Often it is argued that there cannot be an "independent observation language" – that is, a language for the description of observations independently of any prior theoretical framework. (On these two points, see for example [Hesse].) Persons familiar with these issues might suspect that the subjective fact language in *From Brain to Cosmos* is intended to be an independent observation language of some kind. If this were the case, then the line of argument in *From Brain to Cosmos* might be seriously weakened.

Fortunately, this is not the case. The subjective fact language is not intended to be an independent observation language – and the project of *From Brain to Cosmos* does not depend upon its being one. Of course, if one believes in the possibility of an independent observation language, one might be tempted to think that there is some overlap between the classes of subjective fact statements and of independent observation statements. But one does not have to think this, or even to believe in an independent observation language, to recognize that the subjective fact language can be formulated and used for the purposes described in *From Brain to Cosmos*.

Subjective facts are, at bottom, facts about how things seem – although subjective facts differ in certain respects from conventional facts about how things seem. (See *From Brain to Cosmos*, Chs. 2 & 3.) In *From Brain to Cosmos*, I tried to deduce, or at least render plausible, certain conclusions through the use of subjective facts. For that project, it does not matter *how* the subjective facts got

to be true; it only matters *that* they are true. Even if subjective facts are true only by virtue of theory-laden judgments (or theory-influenced perceptions), one still can use such facts as the basis of arguments, as I do in the book.

Note that subjective facts, unlike many "theory-laden" statements, are not subject to abandonment in the face of empirical evidence. This is the case, not because of anything mysterious about subjective facts, but simply because subjective facts are concerned with how things seem and not with how things really are. (The answer to Question 4 below may help with this point.) However, the *subjective content* of a subjective fact certainly can be theory-laden in a sense. For example, a trained weather observer might look out into the sky and immediately notice a rain cloud. The principal subjective fact involved here is one in which it seems in a particular instance that there is a rain cloud. An observer with no knowledge about weather might not see the rain cloud as a rain cloud, but simply as a dark area in the sky. In this case, the subjective fact is one in which it seems in a particular instance that there is a dark area in the sky. This difference in subjective facts reflects a kind of theory-ladenness, though not the strongest possible kind.

In brief, the subjective fact language is not intended to be an independent observation language, and may well be theory-laden in some respects – but none of this has any bearing on the uses of subjective fact presented in *From Brain to Cosmos*.

Question 4. Isn't the very idea of a subjective fact, or the related idea of how things seem, itself theory-laden? Doesn't this possibility cast doubt upon the project of

From Brain to Cosmos?

Answer: We cannot summarily rule out the possibility that these ideas are theory-laden. But even if they were, subjective fact statements still could be true, and the arguments in *From Brain to Cosmos* still would work.

One might think that if the notion of subjective fact (or of how things seem) turned out to be theory-laden, then the project of *From Brain to Cosmos* would be undermined, because that project would be tied to a particular, and perhaps revisable, theoretical standpoint. The only theoretical revision that plausibly could threaten the notion of subjective fact would be a revision that causes the phrase "It seems that...", or its equivalents, to be abandoned. Presumably this would be a revision in our beliefs about mental phenomena. But even if these beliefs were radically revised (and I will not argue that this should happen), there would be no need to abandon the view that it can *seem* that something is the case. The following argument shows one reason why.

Suppose, for the sake of *reductio ad absurdum*, that we adopted some theory that forced us to deny truth to statements of the form "It seems that P." Then we could simply introduce a new word, say "seems-1," and use it in all the situations where we previously would have used "seems." We could take the statement "It seems-1 that P" to be true if and only if one of *those* situations obtained. Then we could decide to redefine the word "seems" to be an abbreviation of "seems-1." In this manner, we could keep using statements of the form "It seems that P" in the customary way, even without the mind-related beliefs that we now automatically associate with such statements. We could do this even if we originally learned how to use and understand "seem" with the help of the old beliefs about

the mind, and even if the circumstances under which the statements are true are picked out with the help of the outmoded beliefs. (We could just consider the old belief system as a device for picking out situations – in much the same way that a game leads to the picking out of a winner.) Thus, the adoption of the new theory could not interfere with our use of statements of the form "It seems that P." The same argument is applicable to subjective fact statements, which are not quite statements of the form "It seems that P."

Question 5. Is *From Brain to Cosmos* an attempt to found all knowledge upon subjective fact?

Answer: No! The project of *From Brain to Cosmos* is something much more modest: an attempt to find out how much metaphysical knowledge can be built upon a specific domain of facts (facts about how things seem). The most that the book accomplishes by way of foundations is the founding of a limited range of metaphysical knowledge – but even limited results of this sort can be interesting. (And even these limited foundations are not unanalyzable. Those who have read the entire book will know what I mean by this last remark.)

Question 6. Could the language of subjective fact be a private language?

Answer: There is no particular reason to think that the language of subjective fact could be a private language. However, for the purposes of the book, it doesn't really matter if it could. *From Brain to Cosmos* contains arguments in which subjective fact statements, or

generalizations built upon such statements, play important roles. If it turns out that one can't learn to use these statements without being part of a linguistic community, this has no bearing on the truth of the subjective fact statements – and the arguments still will go through. What matters to the book is not *how* subjective fact statements come to be used, but *that* they can be used.

Wittgenstein famously argued against the possibility of a private language [Wittgenstein, pars. 256-271]. One's stand on this question does not bear on the project of *From Brain to Cosmos*.

Question 7. Does the concept of subjective fact depend upon folk psychology?

Answer: This question presupposes that there is such a thing as "folk psychology" as some philosophers of mind understand that term. Here I will not address the large issues surrounding folk psychology, and will not try to summarize the debate about this concept, but will focus on the question at hand.

The answer to this question is implicit in the answer to Question 4. Even if our beliefs about the mind were infested with folk psychology and needed to be abandoned or radically revised, there still would be no reason to stop talking about how things seem – and we still could regard subjective fact statements as true or false. (Perhaps the "situations" mentioned in my answer to Question 4 could be neurophysiological or behavioral situations.) As long as some subjective fact statements can be regarded as true, we can use these statements as premises and can argue about them, as I did in *From Brain to Cosmos*. What matters isn't *how* subjective fact statements come to be

true, but *that* they come to be true.

For the record, I am not endorsing the view that folk psychology, as usually understood, really exists – or that it deserves all the attention it has gotten in the literature. I am only pointing out that the idea of folk psychology cannot be used to mount a successful critique of the notion of subjective fact.

References

[Hesse] Mary Hesse, "Is There an Independent Observation Language?" In *The Nature and Function of Scientific Theories*, ed. Robert G. Colodny (Pittsburgh: Univ. of Pittsburgh Press, 1970).

[Wittgenstein] Ludwig Wittgenstein, *Philosophische Untersuchungen / Philosophical Investigations*, trans. G.E.M. Anscombe, 2nd ed. (Oxford: Basil Blackwell, 1958).

Appendix B

Preface

A Meeting of the Methods

This book is a study of a centuries-old philosophical problem and of some philosophical topics on which the study of this problem may shed light. I will not state the book's central problem here, since I will do this in detail in the first chapter. Instead, I will use this preface to make some general remarks about the philosophical methods which I use in the book, and about my motives for using these methods.

Philosophically literate readers will notice that the content of this book differs in some respects from the kind of philosophy which normally is done at American universities. This difference is entirely intentional. The problem which I discuss in this book is a large problem, in the sense that facts and ideas from many different branches of knowledge have significant bearings upon it. This makes the problem difficult to approach through reasoning of the narrowly focused sort which one usually finds in philosophical journal papers. Instead, a more "generalist"

From Brain to Cosmos

approach is needed, drawing on facts from several different fields of philosophy and of the biological and physical sciences. (I have myself published some narrowly focused journal papers, and I found that the writing of this book required a conceptual shift of gears.)

Another difference between the contents of this book and most current academic philosophy has to do with the ways in which the tools of logic and linguistic analysis are deployed. My principal aim is not to analyze concepts or meanings, although I will do that often enough. Instead, my aim is to trace out the logical consequences of a particular set of facts. Whether this approach to the book's problem is successful — or is at least more promising than a strictly analytic approach — is a question which I will leave to readers of the entire book.

From one point of view, my undertaking in this book is quite modest. At no point will I pretend to have a conclusive solution to the problem which is the book's principal theme. Instead, I simply will propose and defend partial and tentative solutions to some of the questions which arise from the study of the book's central problem. Toward the end of the book, I will argue that these solutions point to a new general view of the nature of reality. This general view is not offered as a closed or final "system," but rather as a perspective which may have certain logical advantages over other metaphysical standpoints.

The construction of general metaphysical viewpoints of this sort is not a favored activity in academic philosophy

today. Strictly speaking, this activity goes beyond the scope of analytic philosophy proper, and belongs to the realm of *speculative* philosophy — the type of philosophy which attempts to understand the world in a comprehensive and general fashion. Most of the better-known pre-twentieth-century philosophers, such as Plato, Spinoza, and Descartes, were speculative thinkers. However, speculative philosophy has a bad reputation today, because the speculative systems of philosophy past are felt to be extravagant and defective. Because of this, I should say a few words about my reasons for offering a viewpoint which could be regarded as speculative.

In my opinion, theories in philosophy should serve purposes analogous to those of theories in science. In the physical sciences, a theory is considered successful by virtue of what it is able to explain or describe. The value of a theory is not established solely through the analysis of concepts or through appeals to so-called "standard intuitions." (Indeed, such intuitions often have to yield to the onslaught of stranger but more rational forms of thought, as happened during the birth of quantum mechanics and of relativity.) What is more, a theory need not be free of "loose ends" to be of value to the scientist. Many widely used theories in physics and chemistry are approximate in character, or even contain important logical defects. Such theories still contribute greatly to people's understanding of physical phenomena, although strictly speaking such theories are not quite right. The fact that such a theory describes or explains a wide range of phenomena gives us

confidence that although the theory is not correct, something much like it probably is right.

For precisely the same reason, a philosophical theory can be of interest even if it is not quite right. The methodology of analytic philosophy tends to foster the feeling that a theory must be logically unassailable before it can be acceptable in any sense. A theory about which a bothersome logical question can be raised is felt to be a worthless or even a disreputable theory. If physical scientists had taken this stance, the well-known achievements of their field would remain pipe dreams. In philosophy as in physics, an imperfect theory which lets us understand what we did not understand before may be an interesting and useful approximation to the "right" theory, if there is such a thing as the "right" theory.

If one regards a philosophical theory as a complex logical toy which is absolutely unacceptable unless it is airtight, then the great speculative systems of the past are indeed unacceptable. However, if one regards a theory as a provisional model of reality which may help us to understand something in spite of its defects, then the case against the historic speculative thinkers becomes much weaker. I believe that the latter verdict is the just one. Although speculative philosophy which pretends to be conclusive may be improbable in the face of modern logic, the idea of a speculative philosophy — of a general, overall view of reality — is by no means dead.

Ultimately, the question of the possibility of a logically satisfactory speculative philosophy is an empirical question.

From Brain to Cosmos

We can find out whether such a philosophy is possible only by trying to construct one.

The general view of reality which I have threatened to offer in this book is not a speculative theory of the sort put forward by the thinkers just named. It is a much more tentative, modest, open-ended, and unsystematic endeavor. However, it does address the question of the general nature of reality. In this sense, it is a speculative endeavor — and the methodology which gives rise to it may be regarded as a result of the meeting of modern methods of analysis with the synthetic method of speculative philosophy, which perhaps has not yet outlived its usefulness.

M.F.S.
Los Angeles, California
2000

Appendix C

How Subjective Fact Ties Language to Reality

Mark F. Sharlow

URL: <http://www.eskimo.com/~msharlow>

ABSTRACT

In this note, I point out some implications of the experiential principle* for the nature of the relationship between language and the world. I argue that this principle implies the existence of a certain relationship between linguistic tokens and facts, and that this relationship undermines most critiques of the referentiality of language.

* Sharlow, M.F. *From Brain to Cosmos*. Parkland, FL: Universal Publishers, 2001.

The central metaphysical argument in *From Brain to Cosmos* [1] rests largely upon a thesis called the *experiential principle*. In the book, this principle is stated as follows [1, p. 286]:

Let P be a statement. Suppose that it is logically possible that some subject knows the truth value of P. Then the truth value of P can be determined, to the extent to which it can be known, from facts about what is the case for consciousness events.

In this note, I wish to point out some implications of the experiential principle (hereafter called EP) for the relationship between language and the world. In particular, I will argue that EP, though framed in terms of subjective fact, actually implies the existence of a firm tie between language and *objective* fact -- a tie whose existence makes the reality of an objective world, and the possibility of describing that world, immune to critiques of the referentiality of language. The argument presented here can be regarded as a further development of the informal argument about relativism presented in [1, pp. 309-313].

This essay makes heavy use of two technical notions developed in *From Brain to Cosmos*: the notions of *subjective fact* and *instance of seeming*. Readers who are not familiar with these notions should refer to the book itself [1] and also to reference [2].

Consider a scenario in which someone sees an object -- say, a cardboard box. In this situation, the following subjective fact obtains:

the fact that in some instance of seeming, it seems that there is a box

Call this subjective fact P.

Now suppose that this observer of the box, upon seeing the box, exclaims "There is a box." According to EP, the truth value of this utterance "can be determined, to the extent to which it can be known, from facts about what is the case for consciousness events." In other words, the truth value of this utterance of "There is a box" is determined by certain conditions involving subjective facts, provided that this truth value is knowable. We can think of these conditions as truth conditions for this utterance of "There is a box." We will call them the subjective factual truth conditions for the utterance. (The existence of these conditions does not imply that these truth conditions are the most natural or useful truth conditions for the utterance. One might be able to write a different set of truth conditions that give the utterance the same truth values, but that are more useful for scientific or other purposes.)

The precise formulation of these subjective factual truth conditions is not necessary for our present purposes. What is important is that EP implies that such conditions exist. However, it is safe to assume that if someone actually perceives the box, then the subjective facts associated with that perception are among those which enter into the truth conditions for "There is a box." We must assume that the subjective factual truth conditions can be formulated in a way that involves these subjective facts. Otherwise, we get an unacceptable consequence: we find that the fact that the box is observed is irrelevant to our determining that the box exists. (Certainly, a cardboard box is the kind of a thing whose existence can be inferred from its being observed under suitable conditions.) Thus, the subjective fact P, described above, enters into the subjective factual truth conditions for the utterance of "There is a box."

This connection between P and the truth of the sentence "There is a box" is a connection of a surprisingly deep sort. It is not simply a matter of P playing a role in the observer's knowledge about the box. Instead, the subjective fact P is among the subjective facts that determine, via subjective factual truth conditions, the truth value of the utterance "There is a box." One even can think of P as contributing to the existence of the box -- not, of course, by being a cause of the existence of the box, but in a purely logical way, by playing a role in the subjective factual conditions for the box's existence (see [1], chap. 13).

Now let us modify this scenario. Suppose, as in the original scenario, that a person sees a box, and upon seeing the box, utters the sentence "There is a box." Suppose further that the person perceives the utterance to take place while the box is being seen, and that the person *feels subjectively* that the sentence he is uttering expresses a fact about what he has seen.

This scenario is much like what usually happens when a speaker tries to use language to express facts about objects that the speaker currently is observing. For our present purposes, we won't ask whether the observer in our scenario really is trying to express a fact about the world. We won't try to guess what really goes on inside the observer. Also, we will not yet ask any philosophical questions about whether sentences really can express facts at all. For now, we will only assume that the person *feels* that the sentence expresses the fact that there is a box.

Let us analyze this scenario in terms of subjective facts. To keep the argument to the point, we will temporarily play a Cartesian card and ignore the objective facts about what exists in the scenario: the fact that there is an observer, the fact that there are sentences,

etc. Even if we remain noncommittal as to the existence of observers, sentences, etc., we still can assert that there *seems* to be a sentence. (Indeed, it would be logically possible for there to *seem* to be a language even if there were no real languages at all. Note that this last claim does not depend upon any ideas about nonexistent objects or about the possibility or impossibility of private languages.)

In this scenario, a particular subjective fact obtains: the fact that in a particular instance of seeming, it seems that there is a box. In this same instance of seeming, it also is the case that it seems that there is a sentence -- specifically, the sentence token "There is a box." What is more, in this very same instance of seeming, it *also* is the case that *it seems that the sentence expresses the fact that there is a box*. In this particular instance of seeming, it doesn't just seem that there is a box and also seem that there is a sentence. It also seems that the sentence expresses the fact that the box exists.

Thus, in our new scenario, the following seems to be the case in some instance of seeming:

There is a box, and there is a sentence which expresses that there is a box.

If we wanted to express the corresponding subjective fact in a partially formalized way, we could write:

the fact that in some instance of seeming, it seems to be the case that ((there is a box) and (there is a sentence which expresses that there is a box)).

Call this subjective fact Q.

Now recall our first scenario, in which someone saw a box. In that scenario, there was a subjective fact P. The subjective fact P entered into the subjective factual conditions for the existence of the box. However, one could make the same point about the subjective fact Q in our new scenario. If P can contribute to the existence of the box in the first scenario, then Q can contribute to the existence of the box in the second scenario -- and for the same reason. After all, Q is really just P with some additional content added in.

Note, however, that Q also is a subjective fact in which it seems that a *sentence token* exists -- namely, an utterance of the sentence "There is a box." Thus, Q enters into the conditions for the existence of the sentence token, just as it enters into the conditions for the existence of the box -- and for precisely the same reason. Therefore, *there is an ontological tie between the existence of the sentence and the existence of the box*. The

subjective factual conditions for the existence of the sentence are in part the same as the subjective factual conditions for the existence of the box. These two sets of conditions have at least one subjective fact in common.

Let us say that a sentence *S* *subjectively expresses* a fact *F* if and only if there is an instance of seeming in which it seems that: (a) *S* exists, (b) the fact *F* is the case, and (c) *S* expresses *F*. (Note that for *S* to subjectively express *F*, it only has to *seem* to be the case, in a particular instance, that (a), (b) and (c) are true. The objective truth of (a), (b) and (c) is not required.) The preceding argument shows us that if a sentence token subjectively expresses a fact, then there is a connection, at a fundamental ontological level, between the sentence token and the fact. Therefore, *there is a real tie between language and reality*.

This tie between language and reality arises from experiences that involve *the felt expression of an actual fact that is being experienced*. The preceding argument cannot be made to work for sentences of most other kinds; the argument does not show that every sentence that seems to express a fact, really expresses that fact. The argument shows this only for sentences which are felt to express a fact that actually obtains and that seems to be true. Thus, the argument only shows that sentences of this particular kind have a link to reality. But I think this is much more than many deniers of the referentiality of language should be willing to admit.

A skeptic about the referentiality of language might try to undermine this argument by claiming that there are, strictly speaking, no facts. However, this claim is rebutted implicitly in [1]. According to EP, for any *purported* fact that we care to dream up, if it is logically possible for an observer to know whether that fact obtains, then there are subjective factual conditions such that the fact obtains if and only if those conditions obtain. Thus, given any *purported* fact of this kind, there are subjective factual conditions which, if they obtained, would guarantee that this fact obtains. (If the purported fact is one that does not obtain, then of course these subjective factual conditions do not obtain either.) Hence subjective facts alone can guarantee the reality of other kinds of facts besides subjective facts. (This "guarantee" should be read as logical sufficiency, not as causation.) Once we admit that there are facts which can seem to be the case and seem to be expressed, we must also admit that there can be real ties between language and reality.

In passing, I should mention that this conclusion, and the argument supporting it, have no real connection to Wittgenstein's concept of a private language. (See the discussion of private languages in [2].) The whole issue of private language has surprisingly little to do with the fact that sentences can seem to express facts -- and it is this latter fact that

concerns us here.

The preceding arguments can be adapted to supplement the notion of subjective expression with a notion of *subjective reference*. In our second scenario, there was a tie between the existence of the sentence and the existence of the box. Suppose that instead of just saying "There is a box," the observer gives the box a name: "I'll call that box Henry." To make things more precise and more general, suppose that there is an instance of seeming in which it seems that (a) there is a box, and (b) there is a word, and (c) the word is a name of the box. Then there is a subjective fact in which a word seems to denote an object. This subjective fact contributes to the existence of the object (box) and to the existence of the word as well. It establishes an ontological tie between the object and the word that is supposed to denote the object.

Let us say that a sign of kind S *subjectively denotes* an object of kind F if and only if there is an instance of seeming in which it seems that: (a) a sign of kind S exists, (b) an object of kind F exists, and (c) a sign of kind S denotes an object of kind F. If a word subjectively denotes an object, then there is a connection, at a fundamental ontological level, between the word and the object. The argument for this is similar to our earlier argument about sentences and facts. Again, there is a real tie between language and reality. In case any skeptic tries to get around this by denying that there are objects to be denoted, EP allows us to posit a purported object provided that the subjective factual conditions for the existence of such an object obtain [1, chap. 13]. (Recall our earlier argument about purported facts.)

One can define *subjective reference* in a way analogous to the definition of subjective denotation. Just replace "denotes" with "refers to" in the definiens. Subjective reference and subjective denotation are not the same.

Note that the nature of the chain of events through which an object gets named is irrelevant to this argument. This observation is important because it forestalls certain possible attacks on the argument. Someone might try to undermine the argument given here by claiming that language cannot exist without a social basis; that ostensive definition is not possible without lots of background knowledge; or that private languages are impossible. (All of these three claims have been made by various philosophers of language, though not in connection with our present topic.) But none of this has any bearing on the fact that a word can *seem* to be used to name an object. And this is sufficient for our purposes. It does not matter how the word came into use, or how things came to seem that way.

Someone might try to rebut our argument by claiming that subjective expression (or denotation, or reference) really is not a kind of expression (or denotation, or reference) at all. However, this claim, even if it were right, is irrelevant to the final outcome of the argument. For example, if one thinks that subjective reference isn't really reference, one is free to call subjective reference by another name, like "subreference." (Then subreference would be the kind of referring that non-philosophical language actually does!) Our argument for a deep link between language and reality still will go through. The point is that there is a real link between language and reality -- call the components of that link whatever you will.

Subjective fact is able to underpin certain relationships between language and reality. However, these relationships are not the same as those postulated by traditional correspondence accounts of truth. The links between language and world established by subjective fact are not as "external" and incidental as the relationship between a Cartesian observer and that observer's external world. Instead, these links emerge from the substratum of subjective fact that logically conditions the existence of language and external world alike. It is likely that no critique of the referentiality of language can successfully undermine these relationships, since in a world based on subjective fact, language and reality are not entirely separate to begin with.

One final rebuttal to this argument arises from its apparent circularity. To show that there is a connection between language and reality, we used an argument that speaks of words, sentences, and extralinguistic objects as though we could talk about such things. How can such an argument be used to establish a link between language and world, when the argument apparently presupposes this link? The definitive reply to this criticism makes use of the following observation: the fact that there *seems* to be a language does not imply that there really is a language. If there seems to be a language, then expression, denotation and reference (including denotation of and reference to signs) can seem to occur, whether or not they actually occur. The arguments for EP given in [1] could in principle be formulated even by an observer in a Cartesian dream world. One does not need to presuppose real reference to formulate these arguments; one only needs to presuppose *apparent* reference, and the *appearance* that there is a language. Thus, one can, in principle, formulate in a purely apparent language the argument for the existence of subjective factual conditions for the existence of objects. Using that argument, one can argue (as I have done here) that there is a real tie between language and reality.

Thus, it is possible, given only the ways things seem, to reason to the conclusion that language refers to the world and expresses truths about the world. In this way, subjective fact transcends itself.

References

- [1] Sharlow, Mark F. *From Brain to Cosmos*. Parkland, FL: Universal Publishers, 2001.
- [2] Sharlow, Mark F. "Notes on *From Brain to Cosmos*: Questions and Answers on Subjective Fact," <http://www.eskimo.com/~msharlow/fbtcnote.pdf>.

"How Subjective Fact Ties Language to Reality"
Copyright © 2005 Mark F. Sharlow. Updated 2009. Rebuilt and reformatted 2013.