INFORMATION TO USERS

While the most advanced technology has been used to photograph and reproduce this manuscript, the quality of the reproduction is heavily dependent upon the quality of the material submitted. For example:

- Manuscript pages may have indistinct print. In such cases, the best available copy has been filmed.
- Manuscripts may not always be complete. In such cases, a note will indicate that it is not possible to obtain missing pages.
- Copyrighted material may have been removed from the manuscript. In such cases, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, and charts) are photographed by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each oversize page is also filmed as one exposure and is available, for an additional charge, as a standard 35mm slide or as a 17"x 23" black and white photographic print.

Most photographs reproduce acceptably on positive microfilm or microfiche but lack the clarity on xerographic copies made from the microfilm. For an additional charge, 35mm slides of 6"x 9" black and white photographic prints are available for any photographs or illustrations that cannot be reproduced satisfactorily by xerography.

Order Number 8726826

Epistemic justification and psychological realism

Taylor, James Edward, Ph.D.

The University of Arizona, 1987

Copyright ©1987 by Taylor, James Edward. All rights reserved.

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106

EPISTEMIC JUSTIFICATION AND PSYCHOLOGICAL REALISM

by

James Edward Taylor

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

1 9 8 7

Copyright 1987 James Edward Taylor

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Final Examination Committee, we certify that we have read
the dissertation prepared by James Edward Taylor
entitled Epistemic Justification and Psychological Realism

and recommend that it be accepted as fulfilling the dissertation requirement
for the Degree of Doctor of Philosophy.

Alvin Goldman

8/7/87
Date

John R. Solloch

8/7/87
Date

Kathy Nelson

8/7/87
Date

Date

Date

Final approval and acceptance of this dissertation is contingent upon the
candidate's submission of the final copy of the dissertation to the Graduate
College.

I hereby certify that I have read this dissertation prepared under my
direction and recommend that it be accepted as fulfilling the dissertation
requirement.

Alvin Goldman

Dissertation Director

8/7/87
Date

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the copyright holder.

SIGNED: James E. Taylor

FOR JENNIFER, SARAH AND BENJAMIN

ACKNOWLEDGMENTS

Thanks are due to the following people for contributing to the completion of this dissertation. Professor Robert Wennberg of Westmont College sparked my interest in philosophy with his colorful lectures. I acquired my fascination with epistemology from Professor Stanley Obitts, also of Westmont.

The members of my dissertation committee, Professor Alvin I. Goldman, Professor Keith Lehrer and Professor John L. Pollock, provided me with challenging instruction in epistemology throughout my graduate program. Professor Goldman read every draft of the dissertation and offered many valuable suggestions for its improvement. In addition, I learned a substantial amount of cognitive psychology from Professor Goldman and Professor Robert M. Harnish. Professor Goldman's research concerning the connections between epistemology and psychology provided the impetus for my selection of this issue as a dissertation topic. Professor Stephen Schiffer gave me a solid grounding in philosophy of mind and philosophy of language.

Jody Kraus and Thomas Senor spent numerous hours discussing various parts of the dissertation with me. A conversation with Thomas Senor resulted in my coming to grips with the central idea of this project. Jody Kraus

convinced me that the main argument was worth defending at a time when I was sure it was wrong-headed.

Job interviews at the 1986 American Philosophical Association Eastern Division Meeting provided an opportunity for me to try my ideas out on a random sample of members of the philosophical community. Discussions were particularly lively with the interviewers from Syracuse University, Wellesley College, Occidental College and Bowling Green State University. Campus interviews at Bowling Green State University and Santa Clara University afforded me my first opportunity to read a condensed version of the dissertation and discuss it at length with captive audiences.

Research during the summer of 1987 was made possible by a Summer Research Support Program grant from the University of Arizona Graduate College.

Thanks are due to my parents, Robert and Virginia Taylor, for helping my family and me endure the stresses of finishing a dissertation, finding a job, selling a house, moving and having a baby. The biggest sacrifice was made and the most support has come from Jennifer, my wife. Without her constant encouragement and devotion, this dissertation would not have been written. My daughter, Sarah, expressed her concern by occasionally asking me "How's your dissertation, Daddy?". Benjamin, my son, provided the most joyful interruption of the work by being born. I dedicate this dissertation to these last three people.

TABLE OF CONTENTS

	Page
ABSTRACT.....	ix
1. EPISTEMIC JUSTIFICATION: QUESTIONS AND META-QUESTIONS..	1
Epistemic Justification.....	3
Epistemic Justification as a Species of Justification.....	3
Epistemic Justification and Truth.....	6
Epistemic Justification and Knowledge.....	8
Being Justified vs. Justifying.....	14
Actual vs. Counterfactual Justification.....	15
Summary.....	16
Four Important Questions About Epistemic Justification.....	17
What is the Structure of Epistemic Justification?.....	18
What is the Correct Set of Epistemic Principles?.....	21
What is the Nature of Epistemic Justification?.....	26
Are Humans Capable of Having Justified Beliefs?.....	29
Summary.....	33
Meta-Theory, Epistemic Justification, and Psychology.....	34
The Traditional Thesis.....	37
The Replacement Thesis.....	40
The Dependency Thesis.....	44
Summary.....	47
2. PSYCHOLOGICAL REALISM AND PSYCHOLOGICAL INVESTIGATION.....	48
Stage One: The Constraint of Psychological Realism.....	50
The Denial of Skepticism About Justification.....	51
The Constraints of Realism and Psychological Realism.....	55
Epistemic Justification and Psychological Capacities.....	57
An Objection.....	62
Summary.....	66

TABLE OF CONTENTS--Continued

	Page
Stage Two: The Need for Psychological	
Investigation.....	67
The Indispensability of Psychology.....	68
Three Objections to the Indispensability	
Argument.....	69
Summary.....	83
Conclusion.....	84
3. APPLYING THE CONSTRAINT OF PSYCHOLOGICAL REALISM:	
NON-SCIENTIFIC PSYCHOLOGY.....	87
Overview of the Illustrations.....	88
What the Examples Illustrate About the	
Constraint.....	90
The General Form of the Arguments.....	106
Summary.....	107
Two Appeals to Introspection.....	108
Alston on Chisholm.....	108
Cohen on Lehrer.....	122
Summary.....	138
4. APPLYING THE CONSTRAINT OF PSYCHOLOGICAL REALISM:	
SCIENTIFIC PSYCHOLOGY.....	140
Two Appeals to Cognitive Psychology.....	141
Cherniak on Quine.....	141
Kornblith on Cartesian Epistemology.....	162
Consequences of Widespread Psychological	
Unreality.....	192
The Importance of the Constraint of	
Psychological Realism.....	194
The Indispensability Thesis.....	195
Motivation for My Proposal.....	197
Summary.....	198
5. TOWARD A PSYCHOLOGICALLY REALISTIC ACCOUNT OF	
EPISTEMIC JUSTIFICATION.....	200
An Account-Schema Which Satisfies Both	
Constraints.....	202
Reliabilist Accounts of Epistemic	
Justification.....	215
Is Reliabilism Psychologically Realistic?....	216
Does Reliabilism Satisfy the Intuitiveness	
Constraint?.....	227
Design Accounts of Epistemic Justification.....	232
Problems for the Design Approach.....	234

TABLE OF CONTENTS--Continued

	Page
A Reliabilist Design Account.....	240
Proper Functioning Accounts.....	246
Summary, Conclusion and Prospectus.....	254
APPENDIX: HARMAN ON BAYESIAN EPISTEMOLOGY.....	259
REFERENCES.....	262

ABSTRACT

The central thesis of this dissertation is that it is not possible to determine the nature of epistemic justification apart from psychological investigation. Two sub-theses provide the primary support for this claim. The first sub-thesis is that no account of epistemic justification is correct which requires for the possession of justified beliefs a psychological capacity which humans do not have. A different way of stating this view is that the correct account of epistemic justification must be psychologically realistic. The second sub-thesis is that it is not possible to determine whether an account of epistemic justification is psychologically realistic apart from psychological investigation. In sum, there is a meta-theoretical constraint of psychological realism on accounts of epistemic justification which requires appeal to psychological investigation in its employment.

After defending these proposals, I illustrate how the constraint of psychological realism has been and can be used both to test candidate accounts of epistemic justification and to guide the construction of such an account which is intuitive and psychologically realistic. These two kinds of applications of the constraint can involve either scientific or non-scientific psychological investigation. I give

examples from current epistemological literature of critical employments of the constraint which appeal to both of these kinds of psychological investigation. Finally, in illustrating the role of the constraint of psychological realism in guiding the construction of an account of epistemic justification, I consider both reliabilist views and a variety of positions which feature the notion of cognitive design. I suggest that this latter approach holds out promise for yielding an account of epistemic justification which is both psychologically realistic and intuitive.

CHAPTER ONE

EPISTEMIC JUSTIFICATION: QUESTIONS AND META-QUESTIONS

Two major concerns of recent epistemology are to specify the relations between psychology and epistemology and to investigate the nature of epistemic justification. In "Epistemology Naturalized" (1969), Quine pioneered an attempt to exhibit interesting links between psychology and epistemology. His efforts were followed by those of Harman (1973); Goldman (1978), (1985), and (1986); Kornblith (1982), (1985a), and (1985b); and Cherniak (1986); among others. Of these, Goldman's Epistemology and Cognition (1986) is probably the most comprehensive, large-scale work to date which examines in detail the bearing of psychological investigation on traditional epistemological questions. The second concern, having to do with the nature of epistemic justification, is pursued in the following publications: Lehrer (1974), (1986); Chisholm (1977); Goldman (1979), (1986) and (unpublished); Kornblith (1983); Alston (1985); Pollock (1986); and Plantinga (1986). Many others have studied this question as well.

In this dissertation, I attempt to merge these two projects by defending a thesis about the relationship between determining the correct account of epistemic justification and psychological investigation. More specifi-

cally, I will argue that it is not possible to determine the nature of epistemic justification apart from psychological investigation. To many philosophers, especially those who have been taught to shun "psychologism" in epistemology and logic, this thesis may sound surprising, even preposterous. Nevertheless, the thesis comprises the conclusion of a valid deductive argument with premises which are plausible and relatively uncontroversial. Moreover, this argument does not involve anything as controversial as rejecting either the a priori/a posteriori distinction or the analytic/synthetic distinction. In addition, there is ample evidence in recent epistemological literature that some, if not many philosophers have, at least implicitly, assumed that the thesis is true and have made important use of it. These claims will be defended in chapters two, three and four. However, before proceeding to this defense, it will be helpful to consider some preliminary points. It is to such points that this chapter is devoted.

The rest of this chapter is divided into four sections. First, I will make some clarificatory comments about epistemic justification. Second, I will distinguish four important questions about epistemic justification which are often raised by epistemologists. I will say a little about the relations between these questions and will identify the questions with which we will be most concerned. An important result of this discussion is the recognition of a meta-

theoretical constraint on the answers to some of these questions which can be generated by a certain way of answering another of them. Consideration of this constraint leads to a discussion of some meta-theoretical questions about epistemic justification. This section is the most important one of this chapter, since the primary thesis of this dissertation is meta-theoretical. In this section I will sketch a three-fold taxonomy of positions about the relationship between epistemological questions and psychological investigation. My thesis falls into one of these three categories. In the course of developing this taxonomy, I will discuss a traditional meta-theoretical constraint on accounts of epistemic justification. If my thesis is correct, this constraint will require supplementation by an additional constraint, as we shall see. The last section will contain a summary of the previous three.

Epistemic Justification

Epistemic Justification as a Species of Justification

Justification simpliciter is a generic notion (Bonjour 1985, pp. 5-6). This notion may be characterized generally and roughly as a relation between a cognitive state, event, process or product (or a person's possession of one of these) and an evaluative standard. We can put this characterization schematically as follows in the form of a sufficient condition: A cognitive item or product c is justi-

fied (or a person is justified in possessing or producing this cognitive item or product) if c bears relation R to evaluative standard S. Specific kinds of justification can be generated by means of appropriate replacements of the schematic letters in this schema. For instance, a judge's decision may be justified relative to some legal standard. We might say in such an event that the judge's decision is legally justified. Also, a person's action may be justified in virtue of its satisfying a moral standard. In such a circumstance, we would say that the person's action is morally justified or that she is morally justified in performing that action. Furthermore, prudential justification is usually taken to be a relation between a person's decision or action and a standard of prudential rationality.

It is often assumed that beliefs, in addition to actions and decisions, can be morally justified or prudentially justified (Cohen 1984, p. 279; Moser 1985, p. 1; Pollock 1986, pp. 7-8). For example, Cohen (1984, p. 279) suggests that "the belief of a defense attorney who, in order to provide a better defense, convinces himself that his client is innocent, may be said to be justified (either prudentially or morally)." Moser (1985, p. 1) provides an illustration of a prudentially justified belief. A person washed asea quite far from land may be prudentially justified in believing that he is capable of swimming to shore even if he has overwhelming evidence to the effect that he

is not capable of doing this. Now although these examples are problematic, since they seem to require self-deception and the direct voluntary control of belief (a questionable psychological thesis which will be examined in chapter three), they do make clear that beliefs can be evaluated from different possible points of view. Accordingly, when we say that a belief is justified, we also need to specify the evaluative standard from the standpoint of which the belief is justified. In epistemology, we are particularly interested in evaluations of beliefs from the epistemic point of view. That is, we are interested in the species of justification which is called "epistemic justification."

What is the epistemic point of view? A first step in answering this question is to observe that prudential justification and moral justification are not the same as epistemic justification. The above examples make this clear. The lawyer in Cohen's example may have ample evidence which indicates that his client is guilty. Given this, although he may be morally or prudentially justified in believing his client to be innocent, he would not be epistemically justified in holding this belief. The same is true of Moser's swimmer. The entailments do not hold in the opposite direction either. Pollock (1986, p. 8) asks us to "consider Helen, who has overwhelming evidence that her father is Jack the Ripper." Further, "it may be that if she admitted this to herself it would be psychologically crush-

ing." If Helen believed that her father was Jack the Ripper, this belief of hers would be epistemically justified, but it would not be prudentially justified. It is also clear that epistemically justified beliefs are not necessarily morally justified beliefs. A person may have excellent reasons for thinking ill of another person and yet, assuming that beliefs can be objects of moral evaluation, she might not be morally justified in thinking this if she had promised the person that she would not think ill of him (Pollock 1986, p. 8). Given these considerations, it is clear that epistemic justification is not the same as prudential or moral justification. Furthermore, it is not the case that epistemic justification entails or is entailed by either of these latter kinds of justification.

Epistemic Justification and Truth

What distinguishes epistemic justification from other kinds of justification such as moral and prudential justification? When it comes to spelling out just what is involved in the "epistemic standpoint" which distinguishes epistemic justification from the other types of justification, controversial issues arise. Although it might seem that what makes epistemic justification distinctive is its connection to truth, this connection has been notoriously hard to specify in a non-trivial way (see Cohen, 1984 and Lehrer and Cohen, 1983). Alston, in (1985), says that the

epistemic point of view ". . . is defined by the aim at maximizing truth and minimizing falsity in a large body of beliefs" (p. 59). He includes this in a section of theses concerning justification which he regards as "common ground." In light of Cohen's arguments in (1984), this assumption is contentious. If Alston has something like Cohen's trivial connection to truth in mind, that being justified in believing that p is to be justified in believing that p is true (since one believes that p if and only if one believes that p is true), then his characterization of the epistemic point of view does not distinguish it from non-epistemic points of view. If Alston has something more substantive in mind, which he presumably does, Cohen's arguments show that it is far from uncontroversial. Goldman, in (1986, p. 98), calls the view that Alston propounds "Verific Consequentialism." He distinguishes this view from four other candidate epistemic consequentialisms and argues for it in favor of the others. This is further reason to view the claim as controversial. One point about the truth-connection which seems to have gained virtual universal acceptance is that epistemic justification does not entail truth. If there is a connection between epistemic justification and truth, it must be a connection which leaves epistemic justification fallible. In chapter five, we will discuss the relation between justification and truth again when we examine some reliabilist views of justification.

Epistemic Justification and Knowledge

There is an alternative way of explaining, at least partially, what makes epistemic justification different from other species of justification. This is to say that one thing that makes epistemic justification distinctive is that it is the kind of justification which is required for propositional knowledge. Although this claim is not as controversial as attempts to distinguish epistemic justification in terms of a non-trivial connection with truth, it is contentious nonetheless. Of course, if 'epistemic justification' were defined as "the kind of justification required for propositional knowledge," then the claim would be trivial. But no plausible attempt to characterize the nature of epistemic justification defines it in terms of knowledge. Given that this thesis is not trivial, is it true? It is generally agreed among epistemologists that the concept of propositional knowledge is to be analyzed as true belief plus something that makes true belief knowledge. What this additional component is has been the subject of much dispute. Some theorists contend that this knowledge-making condition includes justification, whereas others deny this. I will side with the former school of thought.

Traditionally, knowledge was thought to consist solely in justified true belief. However, Edmund Gettier's (1963) counterexamples showed that justified true belief is not sufficient for knowledge. This result led to many

attempts to discover what further condition must be added to justified true belief in order to arrive at an analysis of knowledge not subject to Gettier-type counterexamples. Thomas Senor has pointed out in (unpublished) that if 'epistemic justification' had been merely a technical term used to refer to whatever it is that makes true belief knowledge, the Gettier examples would have had no point. It could not possibly be the case that whatever makes true belief knowledge does not make true belief knowledge. If Gettier's examples had been directed against only certain stipulative definitions of such a technical term, the appropriate moral to draw would have been that the correct stipulative definition for the technical term 'epistemic justification' had not yet been found. Instead, of course, philosophers concluded that the Gettier cases showed that epistemic justification is not sufficient to make true belief knowledge. This inference presupposes that epistemic justification is a fixed notion which needs to be supplemented by some further concept or concepts to complete the analysis of knowledge. This is further reason to think that saying that one thing which makes epistemic justification distinctive is that it is necessary for knowledge is not to say something trivial. One thing that both the traditional analysis and attempts to get around Gettier-type counterexamples have in common is the assumption that knowledge requires justification. Gettier-type counterexamples are specified in terms of justifi-

cation. They are examples of cases in which having a justified true belief is not enough for knowing. If one denies that knowledge requires justification, one does not need to worry about Gettier-type cases.

In opposition to the traditional view, and with indifference to Gettier-type counterexamples, some epistemologists have proposed accounts of knowledge in which a condition of justification is lacking. For instance, Dretske (1981, p. 86) characterizes knowledge as "information-caused belief." There are many other variations on the knowledge-without-justification theme (see for instance Nozick 1981). If these non-traditionalists are right, then we cannot distinguish epistemic justification from other types of justification in terms of its being a necessary condition for propositional knowledge. Furthermore, the non-traditional view of knowledge may be employed to generate an objection to the overall thrust of this dissertation. Some who exclude justification from the analysis of knowledge may object to the project developed in this dissertation on the following grounds. They might hold that only knowledge and components of knowledge are legitimate epistemological concerns. Given this and the denial of the thesis that justification is necessary for knowledge, it follows that justification is not a legitimate epistemological concern. Furthermore, it might be argued that if justification is not a proper topic for epistemological inquiry, then it

is not clear why it should receive any philosophical attention at all. I think there is ample reason to doubt the first premise of this argument. There are notions in the literature which, although they are not considered to be components of knowledge, are nevertheless discussed by epistemologists as genuine epistemological subjects. For instance, beliefs are often described as being rational or irrational. Rationality has not traditionally been thought to be a component of knowledge. However, (credal) rationality is arguably a genuine epistemological concern. So the first premise is questionable. In addition, I think it is reasonable to reject the second premise which is the claim that knowledge does not require justification. So even if the first premise were true, the objection would not succeed. Moreover, if the second premise is false, then there is no apparent obstacle to the view that one thing which makes epistemic justification distinctive is its connection with knowledge. Let us consider an attempt to make the second premise of the argument intuitively acceptable.

In a recent paper, William Alston (unpublished, pp. 9-11) argues that knowledge does not require justification. Although Alston does not advance a fully-formulated analysis of knowledge in this paper, a clue to what he has in mind can be found in the following passage:

Now I come to what I take to be the most decisive reason for denying that (justification) is necessary for knowledge. There are, or could be, cases in

which S is receiving information that p in a way that is paradigmatically reliable, tightly under the control of the facts believed, but in which S's total perspective on the world fails to give adequate support to the belief that p (p. 9).

He offers the following as an example of what he takes to be an intuitive case of unjustified true belief which is nonetheless a case of knowledge. S possesses very compelling reasons for thinking that his sensory experience is unreliable. Everyone S knows has conspired to convince him that for some time he has been, and will continue to be, subject to a neurophysiological experiment in which every other day his sensory experiences are artificially produced in the laboratory and are consequently false. In addition, they tell him that the scientists have arranged both that he have no memory of transit to and from the lab and that he experience both real and artificial sensory experiences as a coherent pattern. They provide S with very convincing evidence and he comes to believe justifiably that he can never tell whether or not he is in the laboratory and hence that he is never justified in trusting the deliverances of his senses. One day S ventures into the street and seems to see a truck approaching. As a matter of fact, his perceptual processes are functioning normally and a truck is coming down the street. He momentarily forgets his skepticism, forms a perceptual belief that a truck is coming down the street and steps back. Alston argues that in spite of S's not being justified in believing that a truck is coming down the

street, S knows that a truck is coming down the street since his perceptual processes are functioning normally and reliably, and having momentarily forgotten his doubt, he is certain that a truck is coming.

Cases such as this can be adjudicated only on the basis of intuition. Is it intuitive that S knows that a truck is coming down the street even though he is not justified in believing this? I find it more intuitive to say that S does not know that a truck is coming down the street, precisely because he is not justified in believing this. By my lights, knowing that p requires not having any reason for thinking either that p is false or that one's reasons for p are not good reasons. In the case of S and the approaching truck, although S has no reason to think that it is false that a truck is coming, he has plenty of reasons to doubt the adequacy of his grounds for believing this. That is, he has good reasons for thinking that his senses are not to be trusted. Consequently, pace Alston, this example does not provide strong intuitive grounds for concluding that knowledge does not require justification. Furthermore, it is quite intuitive that knowledge does require justification. So the objection to our project fails. Moreover, the claim that a distinctive feature of epistemic justification is that it is required for propositional knowledge remains unchallenged.

Before moving on to the next section, which concerns four central questions which epistemologists have been concerned to answer about this kind of justification, there are two more distinctions which need to be made in order to clarify further what is meant by 'epistemic justification.'

Being Justified vs. Justifying

There is an important distinction between the state or condition of being justified in believing that p and the act of justifying one's belief that p. It is clear that the former (being justified) is not the same as the latter (justifying). The justification which is required for knowledge is to be construed in the former sense. Alston has stressed the importance of keeping this distinction in mind, most recently in (1985, p. 58). Harman, in (1973), holds that ". . . reasons for which a person believes something cannot be identified with reasons he would sincerely offer if asked to justify his belief" (p. 27). Furthermore, in addition to its not being the case that being justified is the same thing as justifying, it is clear that one can be justified in believing that p without being capable of providing (orally or mentally) justification for one's belief that p. Goldman, in (1979), argues that there are justified beliefs which are such that the cognizer who has them cannot demonstrate that they are justified (p. 15). Any account of epistemic justification which required that one

either justify one's belief or be capable of justifying one's belief, either to oneself or to others, in order to be justified in holding that belief, would be psychologically unrealistic and would consequently lead to skepticism. This is because it is extremely likely that the majority of the justifications for our beliefs, if not all of them, elude our introspective access and verbal report. In chapter two, I will argue that accounts of epistemic justification which have such requirements for justified belief, and are thus psychologically unrealistic, are false. Hence, we can conclude that being epistemically justified in holding a belief does not require that one demonstrate or be capable of demonstrating the justification one has for that belief.

Actual vs. Counterfactual Justification

A second important distinction is that between actual justifiedness and counterfactual justifiedness. Similar distinctions can be found in Firth (1978); Goldman (1979); Pollock (1979), Audi (1980), and Kornblith (1982). Actual justifiedness is a property of actually-held beliefs which are based on what justifies them in the appropriate way. A precise, illuminating account of this "basing relation" has yet to be formulated (attempts can be found in Pappas 1979b and Swain 1979). S is counterfactually justified in believing that p at t if and only if S possesses a set of justifying conditions J for p at t but either (1) S

does not believe that p at t , or (2) S believes that p at t but p is not based on J at t and p is not based on some other set of justifying reasons J' possessed by S at t . Although both actual and counterfactual justifiedness are important and interesting concepts, actual justifiedness is clearly a more plausible candidate for the justification required for knowledge. It is not sufficient for knowing that p , given that p is true, that one possess good reasons for p if either one does not believe that p or one does believe that p but one's believing that p is not based on any justifying conditions for p which one possesses. See Harman (1973, pp. 24-26) for an argument to this effect. This distinction will come up again in chapters three and five.

Summary

This section has been devoted to a preliminary clarification of the notion of epistemic justification. We have observed that epistemic justification is one species of justification among others. Unlike other forms of justification which might be attributed to beliefs, epistemic justification involves a relation between beliefs and an epistemic evaluative standard. That is, when we speak of a belief being epistemically justified we are evaluating that belief from a distinctively epistemic standpoint. It is not easy to give a non-trivial characterization of what is

involved in the "epistemic standpoint." Attempts to tie epistemic justification to truth in some interesting, non-trivial way are controversial. A less tendentious approach is to distinguish epistemic justification from other kinds of justification in terms of its special connection with propositional knowledge. Unlike other kinds of justification, epistemic justification is required for knowledge. This justification required for knowledge is not to be identified with the process of providing a justification for what one believes and does not require that one provide, or be capable of providing, such a justification. Moreover, epistemic justification is actual, not counterfactual justification. That is, a belief is epistemically justified only if it is actually held by a person and is based on justifying reasons or grounds. Although its relation to knowledge does provide a way of distinguishing epistemic justification from other kinds of justification, many other important questions about epistemic justification remain unanswered. Four of the most important of these questions will be discussed in the next section.

Four Important Questions About Epistemic Justification

Epistemologists have been concerned to answer various questions about epistemic justification. Four of the most important of these questions are the following:

(1.1) What is the structure of epistemic justification?

(1.2) What are the specific conditions under which human beliefs of various kinds are epistemically justified?

(1.3) What is the general nature of epistemic justification?

(1.4) Are human beings capable of having epistemically justified beliefs?

Unfortunately, answers to (1.1) - (1.3) have often been called "theories of epistemic justification." This has led to a tendency either to confuse these questions with each other or at least not to recognize or acknowledge the distinctions between them. In this section, I would like briefly to characterize each of these four questions and the form which an answer to each must take. I am particularly interested in distinguishing question (1.3) from the others, since the central thesis of this dissertation concerns a meta-theoretical constraint on answers to (1.3). As we shall see in chapter two, this constraint is generated on the basis of how I answer question (1.4).

What is the Structure of Epistemic Justification?

Answers to this question have traditionally taken the form of either foundation theories of epistemic justification or coherence theories of epistemic justification (for

the purposes of this discussion, I will assume that Pollock's "direct realism" theory of epistemic justification in 1986 is a foundation theory). These families of theories have arisen, in part, as attempts to solve what has become known as the "the epistemic regress problem" (see Alston 1976; Bonjour 1985, pp. 17-25; Moser 1985, pp. 23-26). The epistemic regress problem can be generated on the basis of the simple observation that many of our justified beliefs depend for their justification on other beliefs which we have. Typically, for each justified belief possessed by us which depends for its justification on other of our beliefs, there will be more than one belief on which its justification depends. Each of these latter, supporting beliefs may themselves be supported by other beliefs. So typically, for every justified belief which is justified in virtue of its relation to other beliefs (every "mediately" justified belief), there will be more than one "chain of epistemic dependence" starting with the justified belief and consisting of at least one other "supporting" belief. There are only four forms these chains can take: (1) They can consist of an infinite series of medially justified beliefs; (2) They can consist of a series of medially justified beliefs which circles back to the original "starting" belief; (3) They can terminate with an unjustified belief; or (4) They can terminate with an immediately justified belief.

Traditionally, these considerations have been used to construct an argument by elimination which concludes that some foundation theory is correct. According to foundation theories, there are immediately justified beliefs, and every chain of epistemic dependence must terminate with an immediately justified belief in order for the starting belief to be justified. This conclusion has been thought to be required given the alleged impossibility of chains of types (1) - (3) yielding justified beliefs. However, recently many attempts have been made to argue for the dual thesis that there are no immediately justified beliefs and that only chains of type (2) (or something similar) can yield justified beliefs. These views have been labeled "coherence" theories of the structure of epistemic justification. The reason foundation theories and coherence theories as solutions to the regress problem have been labeled theories of the structure of epistemic justification is that they can be construed as answering the question "does knowledge rest on a foundation (the immediately justified beliefs) or is the structure of justification more like a foundationless web or nebula?"

Some observations are in order here before we move on to consider the next important question about epistemic justification. The first is that coherence theories of epistemic justification have not only been offered as answers to question (1.1) concerning the structure of epis-

temic justification (see for instance Bonjour 1985). They have also been offered as answers to question (1.3), that is, as general accounts of epistemic justification (see Lehrer 1974, p. 198). In addition, they have been considered to be candidate theories for specifying the character of epistemic principles, the principles which state the conditions under which beliefs of various kinds are justified (Goldman 1986, pp. 194-198). That is, they have been thought to be potential answers to question (1.2) (or at least they have been thought to specify the form that epistemic principles must take). In the same way, although I am not aware of any attempt to answer question (1.3) concerning the nature of epistemic justification with a foundation theory of justification, foundation theories have been offered as answers to both (1.1) and (1.2) as well. These considerations only serve to emphasize the possibility for confusion engendered by the use of the phrase "theory of epistemic justification" when there is no attempt to specify which question or questions the theory purports to answer.

What is the Correct Set of Epistemic Principles?

The answer to question (1.2) would take the form of a list of epistemic principles. An epistemic principle is a universally quantified conditional in which the antecedent states a condition which is sufficient for the justification

of a particular kind of belief. An example of an epistemic principle can be used for illustration:

- (1.5) For any human S and any object O, if O looks red to S and S has no reason to believe either that it is not the case that O is red or that O's looking red is not a reason to believe that O is red, then S is justified in believing that O is red.

The universal quantification over human beings in this principle represents one way in which sets of epistemic principles can be construed. Restricting the domain in this manner makes epistemic principles species-relative. Given this treatment of epistemic principles, (1.5) and other principles which quantify over humans would apply only to humans. If humans are capable of possessing epistemically justified beliefs (that is, if the answer to question (1.4) is "yes,") then humans can satisfy the antecedents of some or all of these epistemic principles which apply to them. If humans are not capable of possessing epistemically justified beliefs, then although these principles will apply to them, no human could satisfy the antecedent of any of them, and they would be vacuously true. A different set of epistemic principles would apply to other conceivable psychological species. For instance, a non-human species might be incapable of forming beliefs on the basis of vision. However, members of this species may acquire some of their

beliefs on the basis of radar. Thus, a principle analogous to (1.5) might be constructed which quantifies over the members of this latter species and which specifies the conditions under which beliefs acquired on the basis of radar are epistemically justified. Consequently, given this approach to epistemic principles, for each different psychological species, at least some of the epistemic principles which apply to that species will differ in content from the epistemic principles which apply to human beings.

Goldman has called my attention (in conversation) to an alternative way of formulating epistemic principles. According to this method, instead of quantifying over only one psychological species, with the result of there being a different set of epistemic principles for each species, epistemic principles are specified so as to quantify over all psychological species. A consequence of this is that there would be only one set of epistemic principles which applies to all psychological species. On this view, if a given species does not have a psychological capacity which is specified in the antecedent of a particular epistemic principle, then that epistemic principle will be satisfied vacuously by that psychological species. A given species may or may not have the psychological capacities required for the possession of epistemically justified beliefs. For instance, a species may or may not have the capacity to possess beliefs. If a species lacks the required capacities

for justified beliefs, then, although the one set of epistemic principles will apply to that species, all of the principles would be vacuously true of that species. That is, there would be no epistemic principle which is such that a member of that species can satisfy the antecedent. Of course, no member of such a species would possess epistemically justified beliefs. On the other hand, if a species possesses the capacities required for justified beliefs, then, although the entire set of epistemic principles would apply to that species, only a subset of this set would have antecedents which can be satisfied by members of the species.

Question (1.2) concerns the specific conditions under which human beliefs of various kinds are epistemically justified. On the first approach, described above, the answer to this question would consist in a list of epistemic principles which quantifies over only members of the human species. On the second approach, the answer would take the form of a subset of the one set of epistemic principles which quantifies over all psychological species. On either approach, whether humans satisfy the antecedents of any of the epistemic principles which apply to them depends on whether they are capable of possessing epistemically justified beliefs. This issue will be discussed in the context of considering the answer to question (1.4).

Another feature of (1.5) which should be pointed out is that the second conjunct of the antecedent specifies two alternative ways in which a potential justifying ground could fail to confer justification on a belief. These alternatives represent two ways in which the stated ground could be "defeated" as a source of justification for the visual belief in question. The general way of referring to this constituent of epistemic principles is to say that epistemic justification is "defeasible." I will have a little more to say about defeasibility in chapter five.

We are now in a position to make a general observation about the relation between question (1.1) and question (1.2). One's answer to question (1.1) will determine, at least partially, one's answer to question (1.2). If a form of foundationalism is true, then there will need to be epistemic principles which specify how immediately justified beliefs are justified and other epistemic principles which state the conditions under which mediately justified beliefs are justified. If a version of coherentism is true, then there will not be both of these kinds of epistemic principles. The form that epistemic principles would take given the truth of some kind of coherentism depends on the kind of coherentism involved.

What is the Nature of Epistemic Justification?

The answer to question (1.3) would take the form of a condition which is logically necessary and sufficient for any belief to be epistemically justified. Questions (1.1) and (1.2) can be answered independently of possessing an answer to question (1.3). If one has accepted a given view about the structure of epistemic justification, one may still wonder what general feature is shared by all and only the beliefs which have this structure. Moreover, although it might be thought that our having a complete list of epistemic principles would reveal to us everything there is to know about epistemic justification, such a list would lack the generality and explanatory power needed for an adequate answer to (1.3). In "What is Justified Belief?" (1979), Goldman puts this point as follows:

The aim of this paper is to sketch a theory of justified belief. What I have in mind is an explanatory theory, one that explains in a general way why certain beliefs are counted as justified and others as unjustified Since I seek an explanatory theory, i.e., one that clarifies the underlying source of justificational status, it is not enough for a theory to state 'correct' necessary and sufficient conditions. Its conditions must also be appropriately deep or revelatory. Suppose, for example, that the following sufficient condition of justified belief is offered: 'If S senses redly at t and S believes at t that he is sensing redly, then S's belief at t that he is sensing redly is justified.' This is not the kind of principle I seek; for, even if it is correct, it leaves unexplained why a person who senses redly and believes that he does, believes this justifiably. What is distinctive about the state of sensing redly, or 'phenomenal' states in general? A theory of justified belief of the kind I seek must answer this question,

and hence it must be couched at a suitably deep, general, or abstract level (pp. 1, 2).

Alston's recognition of the insufficiency of knowledge of epistemic principles for knowledge of the nature of epistemic justification is displayed by his remark that "we want to know not only when our beliefs are justified, but also what it is to enjoy that status" (1985, p. 57). Socrates's dissatisfaction with mere examples of justice as answers to his question concerning the nature of justice illustrates his endorsement of an analogous point. A more specific reason for denying that the answer to (1.2) could also be an answer to (1.3) is that, although the answer to (1.2) is relative to the human species, as we have seen, the answer to (1.3) must be species-neutral. It must specify a single general condition which is necessary and sufficient for the justification of any kind of belief of any psychological species. It is quite reasonable to suppose that some possible non-humans could have justified beliefs. In sum, it is possible to have answers to (1.1) and (1.2) and yet not have a complete answer to (1.3).

Although knowing the content of the epistemic principles which apply to human beings would not be sufficient for knowing what the general nature of epistemic justification is, it may be the case that this latter knowledge would suffice for determining a property of any correct set of epistemic principles, human or otherwise. For instance, in

(1986), Goldman defends the following answer to (1.3):

- S's believing p at time t is justified if and only if
- (a) S's believing p at t is permitted by a right system of J-rules (justificational rules), and
 - (b) this permission is not undermined by S's cognitive state at t. (p. 63).

and

A J-rule system R is right if and only if R permits certain (basic) psychological processes, and the instantiation of these processes would result in a truth ratio of beliefs that meets some specified high threshold (greater than .50) (p. 106).

A "J-rule system" is (roughly) a set of epistemic principles of a certain sort. If Goldman's view is right, then a set of epistemic principles is correct only if it has the property specified in the analysis of the second passage quoted. On the other hand, another example demonstrates the possibility that a general account of epistemic justification would not yield a characterization of any correct set of epistemic principles. Pollock defends the following general account of epistemic justification in (1986): "A person's belief is justified if and only if he holds it in conformance to his epistemic norms" (p. 168). For Pollock, epistemic norms are epistemic principles which play a certain kind of regulative role in reasoning. They are ". . . norms that actually govern our reasoning" (p. 168). It is clear that Pollock's account does not characterize correct sets of epistemic norms in terms of a general property. We

will consider these views of Goldman and Pollock in more detail in chapter five.

One reason that it is important to have an answer to question (1.3) is that, given that epistemic justification is a necessary condition of propositional knowledge, we will not have a complete account of what knowledge is without knowing what epistemic justification is. The question concerning the nature of knowledge goes back at least to Plato's Theatatus. Recent attempts to specify the nature of epistemic justification are continuations of this ancient project. This dissertation is an attempt to contribute to the program to understand what knowledge is. The last epistemological question which we shall consider, before moving on to some meta-epistemological questions, is also as old as ancient Greek philosophy. The larger question is whether humans are capable of having knowledge. Given the relationship between knowledge and epistemic justification affirmed above, the answer to this question depends on whether humans are capable of having justified beliefs.

Are Humans Capable of Having Justified Beliefs?

Negative answers to this question are skeptical answers. Positive answers are anti-skeptical answers. There are various kinds of epistemological skepticism. The kind that concerns us here I will label "J-skepticism(H)". This is the view that humans are not capable of possessing

epistemically justified beliefs. The truth of J-skepticism(H) would leave open the possibility that other possible psychological species are capable of possessing justified beliefs. This is why I have indexed the view to human beings. If J-skepticism(H) were true, although humans would not have what it takes to satisfy the requirements for justified beliefs, some other possible species might. On the other hand, the falsity of J-skepticism(H) would leave open the possibility that a skeptical view about the justification of the beliefs of some possible non-human species is true. It would also leave open the possibility of other forms of skepticism which apply to humans, such as skepticism with respect to the human capability of possessing knowledge. This is because there are skeptical arguments concerning knowledge which do not depend on the denial of human possession of or capability of possessing justified beliefs.

What is the relationship between question (1.4) and the other three questions? Consider first question (1.1). I stated that the two rival answers to this question, foundation theories and coherence theories, can be considered to be potential answers to the traditional epistemic regress problem. These two families of theories constitute two out of four alternative characterizations of the chains of epistemic dependence which lead to mediately justified beliefs. One candidate answer to the epistemic regress problem is

that none of these four kinds of chains can possibly transmit epistemic justification. If this is the case, then the answer to question (1.4) is "no". That is, J-skepticism(H) is one alternative answer to the epistemic regress problem. If this were the case, then question (1.1) would rest on a false presupposition. If none of the possible kinds of chains of epistemic dependence can transmit justification, then there is no structure to epistemic justification since no one could have epistemically justified beliefs. There is also a way in which some proposed answers to (1.1) might lead to J-skepticism(H). If a foundation theory or a coherence theory of the structure of epistemic justification should require for the possession of justified beliefs some capacity which human beings do not possess, then such a view would entail that the answer to question (1.4) is negative.

Consider next question (1.2). The answer to this question will specify a set of epistemic principles which apply to human beings. In effect, this set will comprise the evaluative standard according to which human beliefs are assessed to determine whether or not they are epistemically justified. Now there are candidate answers to (1.2) consisting of sets of epistemic principles which are such that humans are not capable of fulfilling the conditions specified in the antecedent of the principles. Given this, the correct set of (human) epistemic principles would not specify the conditions under which human beliefs are justified

since human beliefs would fail to satisfy the requirements for justifiedness. Consequently, the answer to the question "what is the correct set of (human) epistemic principles?" may entail J-skepticism(H). Moreover, question (1.2) would also be based on a false presupposition, since it would not be the case that humans have any justified beliefs.

The final question to consider in terms of its relationship to question (1.4) is question (1.3). The interplay of these two questions will play a crucial role in chapter two. Candidate answers to question (1.3) concerning the nature of epistemic justification may incorporate a requirement for the possession of justified beliefs which human beings cannot satisfy. Such an answer to (1.3) would entail a negative answer to question (1.4).

In sum, some candidate answers to questions (1.1) - (1.3) will entail a negative answer to (1.4). That is, they will entail that humans are not capable of possessing epistemically justified beliefs. This is the view which we have called "J-skepticism(H)." What is important for our purposes is to recognize that, given this, a positive answer to (1.4) will function as a constraint on answers to (1.1) - (1.3). That is, the denial of J-skepticism(H) imposes a constraint on theories of epistemic justification. If J-skepticism(H) is false, then any epistemological theory which entails J-skepticism(H) is false. One of the main purposes of chapter two is to defend a certain form of such

a constraint on answers to question (1.3). In spite of this limitation of focus to a constraint on question (1.3), if the arguments I offer in the next chapter are sound, there is an analogous constraint on the answers to (1.1) and (1.2) as well. Thus, if my case for a constraint on general accounts of epistemic justification succeeds, then there will be a structure of epistemic justification and there will be conditions under which human beliefs are epistemically justified, since my thesis entails that humans are capable of possessing, and actually do possess, epistemically justified beliefs. So any theory is false which entails that there is no such structure or that such conditions do not exist. Although my arguments are generalizable in this way, I will ignore these wider implications in most of the dissertation, since question (1.3) and questions about (1.3) will be more than enough to occupy us.

Summary

In this section I have distinguished four important questions which epistemologists have attempted to answer. One question concerns the structure of epistemic justification and another concerns when our beliefs are epistemically justified. The question which will occupy most of our attention has to do with the nature of epistemic justification. Finally, the last question concerns the capability of human beings to have epistemically justified beliefs. We

have seen that, although these four questions are indeed separate questions, there are interesting and important relations between them. The most important relation for our purposes is the one between question (1.4) and questions (1.1) - (1.3). Some answers to the latter three questions entail a negative response to the former. Given this, a positive answer to (1.4) will entail the falsity of the theories which entail that the answer to (1.4) is negative. So the denial of J-skepticism(H) yields a meta-theoretical constraint on theories of epistemic justification. We will see in chapter two how the application of a form of this meta-theoretical constraint must involve an appeal to psychological investigation. In the next section, we will consider a taxonomy of meta-theoretical views about the relation between psychological investigation and epistemological theorizing. My endorsement of the meta-theoretical constraint generated by the denial of J-skepticism(H) commits me to one of the meta-theoretical views in this taxonomy.

Meta-Theory, Epistemic Justification, and Psychology

In the previous section I stated that the answer to question (1.3) will take the form of a general condition which is logically necessary and sufficient for the epistemic justification of any kind of belief of members of any psychological species. The primary locution which I will use to describe the activity of providing a correct answer to

(1.3) is "giving a correct account of the nature of epistemic justification." Now of course as philosophers we are not only concerned with providing true accounts of notions which interest us but also with providing good reasons for thinking that these accounts are true. Thus, I will not only talk about giving a correct account of the nature of epistemic justification, I will also speak of determining the nature of epistemic justification. This latter activity has an epistemological component which is lacking in the former. One determines the correct account of epistemic justification if and only if (1) one gives the correct account of epistemic justification and (2) one has good reasons for thinking that it is the correct account. Given this, the central question which will concern us in this section is "What does one need to know in order to determine the correct account of epistemic justification?" As I said at the beginning of this chapter, the main thesis of this dissertation is that it is not possible to determine the nature of epistemic justification apart from psychological investigation. According to this thesis, one thing we need to know in order to determine the nature of epistemic justification is some psychological facts about human beings. This meta-theoretical position is contentious. In the rest of this section, I will discuss this view and two other approaches to the relation between psychological investiga-

tion and determining the correct account of epistemic justification.

Before we consider these three meta-epistemological positions, two observations need to be made. In the first place, it should be kept in mind that, although we will be talking exclusively about meta-theoretical positions concerning determining the answer to question (1.3), an analogous taxonomy could be constructed with respect to questions (1.1) and (1.2). So these meta-theoretical positions would apply to any kind of theory of epistemic justification. Secondly, I intend to use the term 'psychological investigation' rather broadly to encompass both scientific and non-scientific methods of psychological investigation. Scientific methods include the methods of experimental cognitive psychology. Non-scientific methods include introspection and informal observation of others' behavior outside of the context of a controlled experiment. I will have more to say about all of these varieties of psychological investigation in chapter three.

With these preliminary clarifications in mind, we can proceed to the meta-theoretical taxonomy. The taxonomy of meta-theoretical views can be generated by considering the following question: "Can one determine the correct answers to important questions about epistemic justification apart from psychological investigation?" Given our special interest in question (1.3) this question becomes:

(1.6) Is it possible to determine the correct account of epistemic justification apart from psychological investigation?

I will call the view which gives an affirmative answer to this question "the traditional thesis." My label for the view which gives a negative answer to (1.6) is "the thesis of epistemological naturalism." This latter view divides into two distinct sub-views. The first of these sub-views I will refer to as "the replacement thesis," and I will call the second sub-view "the dependency thesis" (Kornblith uses some of these labels in a similar but somewhat different taxonomy in 1985b, pp. 1-13). The overall thesis of this dissertation is the dependency thesis. Consequently, it falls under the general rubric of "naturalistic epistemology." We now turn to a detailed examination of each of these three views.

The Traditional Thesis

According to the traditional thesis, it is possible to determine the correct account of epistemic justification apart from psychological investigation. In chapter two, I will state an argument for this thesis. For now, my concern is merely to characterize the view and contrast it with the two naturalistic epistemology theses. According to the traditional thesis, semantic intuitions about epistemic justification are the only considerations to which one must

appeal in the construction and testing of candidate accounts of epistemic justification. Another way of putting this is to say that the only meta-theoretical constraint on accounts of epistemic justification recognized by the traditional view is a constraint of intuitiveness. According to the traditional thesis, satisfaction of this constraint is considered to be both necessary and sufficient for determining the correct account of epistemic justification. Given this, one has determined the correct account of epistemic justification if and only if the account is intuitive.

What is it for an account of epistemic justification to be intuitive? Semantic intuitions about epistemic justification take the form of the critical judgments we would make concerning the justificational status of beliefs. These critical judgments exemplify procedural knowledge about justification. That is, they exemplify our ability to discriminate between justified and unjustified beliefs. The construction of an account of epistemic justification on the basis of such instances of procedural knowledge about the justificational status of beliefs is a type of hypothesis-formation. Candidate accounts of epistemic justification are hypotheses which are subject to test by our intuitions about epistemic justification. If an account of epistemic justification has the consequence that a belief which we would intuitively regard to be unjustified is justified, or vice versa, then the account fails to be intuitive. Any

account which fails to be intuitive can be concluded to be false, given that satisfaction of the intuitiveness constraint is a necessary condition for determining the correct account of epistemic justification. Given the sufficiency of this constraint as a meta-theoretical constraint on such accounts, an intuitive account (one which is not subject to the type of intuitive counterexample described above) would be a correct account.

I am using 'semantic intuitions' in a narrow sense to include only a priori considerations. In chapter two, I will contrast this narrow sense of intuitiveness with a broader sense which includes a posteriori considerations as well. Given this narrow construal of 'intuitiveness', the traditional thesis can be recast as the view that a priori considerations about epistemic justification are necessary and sufficient for determining the correct account of epistemic justification. As we shall see, the replacement thesis holds that narrow intuitiveness is neither necessary nor sufficient for this purpose. On the other hand, the dependency thesis denies that semantic intuitions are sufficient for determining the nature of epistemic justification but does not deny that intuitiveness is necessary for this task. Consequently, one thing that makes a meta-epistemological view a type of naturalistic epistemology, according to this taxonomy, is that it denies the sufficiency of narrow intuitiveness for answering epistemological questions. In

the present case, the epistemological question at issue is (1.3). Let us examine these latter two naturalistic views more closely.

The Replacement Thesis

We have seen that the traditional thesis answers question (1.6) in the affirmative. That is, according to this thesis, it is possible to answer epistemological questions such as (1.3) apart from psychological investigation. This is because, according to the traditional thesis, the sheerly a priori considerations provided by semantic intuitions about epistemic justification are both necessary and sufficient for answering these questions. In particular, these considerations are necessary and sufficient for determining the correct account of epistemic justification. The replacement thesis constitutes the more extreme of the two negative responses to (1.6) which we will consider. This is because it holds that a priori considerations are neither necessary nor sufficient for answering question (1.3). Instead, according to the replacement thesis, a priori considerations about epistemic justification are entirely replaced by a posteriori considerations for the purpose of determining the correct account of epistemic justification.

The classical statement of the replacement thesis was made by W. V. O. Quine in "Epistemology Naturalized":

Epistemology still goes on, though in a new setting and a clarified status. Epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science. It studies a natural phenomenon, viz., a physical human subject. This human subject is accorded a certain experimentally controlled input - certain patterns of irradiation in assorted frequencies, for instance - and in the fullness of time the subject delivers as output a description of the three dimensional external world and its history. The relation between the meager input and the torrential output is a relation that we are prompted to study for somewhat the same reasons that always prompted epistemology; namely, in order to see how evidence relates to theory, and in what ways one's theory of nature transcends any available evidence (1985, pp. 23-24).

Two observations should be made about this passage as a statement of the replacement thesis understood as a view about what we need to know in order to determine the correct answer to question (1.3). First, although it is not clear whether Quine intends to be talking about epistemic justification here, it is reasonable to suppose that his remarks about the relation between evidence and theory indicate that what he has in mind is something quite similar to the relation between epistemically justifying reasons or grounds and beliefs. Second, the brand of psychological investigation suggested by Quine's comments is behavioral psychology. The replacement thesis need not be restricted to this approach. Quine's remarks could be modified in such a way as to describe the replacement thesis in terms of the replacement of traditional a priori epistemological theorizing by any psychological methodology whatsoever. Recall that, in this

dissertation, I will be using 'psychological investigation' to include the methods of scientific cognitive psychology and more informal methods, such as unaided introspection. My usage reflects the general disparagement of behavioral psychology which pervades current philosophical literature.

Now although I will not consider Quine's argument for the thesis contained in the passage quoted above, I would like to make two criticisms of the replacement thesis. The argument which I offer in chapter two for a negative answer to question (1.6) places my view in the naturalistic epistemology category. My endorsement of the following criticisms of the replacement thesis explains my rejection of this latter version of naturalistic epistemology in favor of what I have called the dependency thesis, the view we will examine in the next section. These two objections to the replacement thesis are aimed at the claim that psychological investigation is sufficient as a source of knowledge for determining the correct account of epistemic justification. I agree that psychological investigation is necessary for this purpose.

Why are the a posteriori considerations yielded by psychological investigation not sufficient for determining the correct account of epistemic justification? There are at least two good reasons for this conclusion. The first is that psychological investigation is a purely descriptive enterprise and epistemic justification is at least partially

an evaluative notion. It is doubtful that a partially evaluative notion could be explicated fully in terms of merely descriptive considerations. Thus, it is doubtful that the nature of epistemic justification can be completely determined on the basis of psychological investigation alone. The second reason for rejecting the view that psychological investigation is sufficient for determining the correct answer to question (1.3) is that it is difficult to see how we could even begin to answer this question without making use of some a priori considerations. Traditionally, the questions "What is knowledge?" and "What is epistemic justification?" have been considered to be philosophical questions. Philosophical questions are questions which require at least some appeal to a priori considerations in order to answer them. An answer to these questions couched in purely psychological terms would seem either to miss the point of the question or to involve a change of subject. It seems most reasonable to conclude that it is not possible to determine the answer to (1.3) without some appeal to a priori considerations. Consequently, in view of this and the preceding criticism of the replacement thesis, we can set the replacement thesis aside as a view about the relationship between psychological investigation and determining the correct account of the nature of epistemic justification.

The Dependency Thesis

As a sub-thesis of the naturalistic epistemology thesis, the dependency thesis holds that it is not possible to determine the correct account of epistemic justification apart from psychological investigation. As we have seen, the replacement thesis agrees with this claim in its endorsement of the view that psychological investigation is necessary for determining the correct answer to (1.3). However, the replacement thesis entails the further proposition that psychological investigation is sufficient for determining the nature of epistemic justification as well. We have just seen that there are good reasons to reject this additional claim. Given this, the dependency thesis can be viewed as a mediating position between the other two. It is the view that the a posteriori considerations yielded by psychological investigation are necessary but not sufficient for determining the correct account of epistemic justification. Determining the answer to (1.3) depends, partially, but not entirely, on knowing some psychological facts.

What are the general reasons for my acceptance of the dependency thesis? In the summary section of the previous major section, entitled "Four Important Questions About Epistemic Justification," I stated that my endorsement of a meta-theoretical constraint generated by the denial of J-skepticism(H) commits me to a view about the relation between psychological investigation and determining the cor-

rect account of epistemic justification. In chapter two, I will argue for this meta-theoretical constraint, which I will call "the constraint of psychological realism." I will also provide an argument which purports to show that psychological investigation is required in the employment of this constraint. Consequently, my acceptance of the constraint of psychological realism commits me to the thesis of epistemological naturalism. This view, coupled with the rejection of the replacement thesis, entails the dependency thesis.

Are there any additional meta-theoretical constraints on accounts of epistemic justification? The dependency thesis, as stated, leaves this question unanswered. However, I will supplement the dependency thesis by endorsing a second constraint. One of my criticisms of the replacement thesis is that it denies the necessity to appeal to a priori considerations, such as semantic intuitions about epistemic justification, for the purpose of determining the correct account of epistemic justification. This suggests that the traditional thesis is partially on the right track, since it makes satisfaction of a constraint of intuitiveness a necessary condition for determining the correct answer to (1.3). However, the traditional thesis fails in its insistence that satisfaction of the intuitiveness constraint is sufficient for this purpose. So both the traditional thesis and the replacement thesis founder in virtue of their recognizing only one kind of meta-

theoretical constraint on accounts of epistemic justification. In the case of the traditional thesis this is an a priori constraint. The replacement thesis recognizes only an a posteriori constraint. The way to avoid the problems of these two views is to endorse both an a priori constraint and an a posteriori constraint on accounts of epistemic justification. My endorsement of the anti-skeptical constraint of psychological realism to be defended in the next chapter is an instance of the latter kind of constraint, in virtue of its entailing the thesis of epistemological naturalism. In addition, I will supplement the constraint of psychological realism with a constraint of intuitiveness in order to avoid the problems of the replacement thesis. Satisfaction of these two constraints is necessary for determining the correct account of epistemic justification. Satisfaction of either alone is not sufficient for this purpose. I do not have an argument for the claim that satisfaction of both constraints by a single account is sufficient to show that the account is correct. However, although I will defend a general approach to epistemic justification in chapter five, I will not argue that it is correct. I will make the more modest claim that it is both intuitive and psychologically realistic. Of course, if having both of these properties is sufficient for having good reason to say that an account or account-schema is correct, then, if my proposal is intuitive and psychologically real-

istic, there would be satisfactory grounds for concluding that it is correct.

Summary

In this chapter, I have discussed some preliminary issues concerning epistemic justification, questions about epistemic justification and meta-questions about epistemic justification. Epistemic justification is the species of justification which is required for propositional knowledge. Among the important questions which philosophers have asked about this kind of justification is "what is the nature of epistemic justification?" If humans are capable of possessing epistemically justified beliefs, then there is a constraint on answers to this question about the nature of justification. As we shall see in chapter two, this constraint entails the thesis of epistemological naturalism, the view that psychological investigation is necessary for determining the correct account of epistemic justification. Given the denial of the claim that psychological investigation is also sufficient for this purpose, it follows that the dependency thesis is true. In addition, there is a constraint of intuitiveness on accounts of justification. In this dissertation, I will defend and illustrate the constraint of psychological realism. It is to the defense of this constraint and of the thesis of epistemological naturalism which ultimately follows from it that we now turn.

CHAPTER TWO

PSYCHOLOGICAL REALISM AND PSYCHOLOGICAL INVESTIGATION

Now that we have discussed some questions and meta-questions about epistemic justification, we are in a better position to answer questions (1.6) and (1.3). Recall that question (1.3) is the following: What is the nature of epistemic justification? Before we can answer this question we must first provide an answer to question (1.6). That is, we must determine whether it is possible to answer question (1.3) apart from psychological investigation. (Recall that I am construing 'psychological investigation' broadly to include both scientific psychology and non-scientific forms of psychological investigation such as unaided introspection.) This chapter is devoted to answering question (1.6). I will provide an argument which purports to show that it is not possible to determine the nature of epistemic justification apart from psychological investigation. This, recall, is the thesis of epistemological naturalism. Given our rejection of the replacement thesis in chapter one, this argument supports the dependency thesis. That is, these considerations entail that psychological investigation is necessary but not sufficient for determining the nature of epistemic justification.

Before I develop the defense for the thesis of

epistemological naturalism, I will state a natural line of argument in support of the traditional thesis, which is its denial. This latter argument has quite a bit of appeal. Nevertheless, if the arguments I defend in this chapter are sound, the following argument is unsound: "Giving a general account of the nature of epistemic justification is just to analyze the concept of epistemic justification. Conceptual analysis is an entirely a priori enterprise. As such, there is no need to appeal to any a posteriori considerations, such as those yielded by psychological investigation. Consequently,

(2.1) it is possible to determine the nature of epistemic justification apart from psychological investigation."

The central question of this chapter is whether (2.1) is true. As I said above, I will argue that (2.1) is false.

This conclusion is in direct conflict with stances taken by other epistemologists who endorse links between psychology and epistemology. For instance, Alvin Goldman has argued that, although cognitive psychology is needed in order to determine the content of epistemic principles (i.e., in order to answer question (1.3)), it is not required in the determination of what epistemic justification is (Goldman 1986, p. 66). For Goldman, this latter task is at the "foundational" level of the epistemological enterprise. The contribution which cognitive psychology

makes to epistemology occurs after the conceptual foundations are constructed (1986, pp. 8-9). So, in (1986), Goldman accepts (2.1). The present thesis is that the proper foundation cannot be built by philosophy alone. Psychological investigation (both scientific and nonscientific) is required at that level as well. This is a case in which philosophers cannot leave the empirical questions to the psychologists.

The argument for this claim proceeds in two stages. The first stage is based upon the notion of psychological realism which I will introduce below. The conclusion of the first stage is that the correct account of epistemic justification must be psychologically realistic. In the second stage, I argue that the only way to determine whether a candidate account of epistemic justification is psychologically realistic is to appeal to psychological investigation. If these two arguments are sound, it follows (given a certain additional quite reasonable assumption) that (2.1) is false.

Stage One: The Constraint of Psychological Realism

In this first stage of the argument against (2.1), I will argue, on the basis of a certain anti-skeptical thesis, for what I will call a constraint of realism on accounts of epistemic justification. This constraint entails a constraint of psychological realism on such accounts. Before going on to the second stage of the argument against (2.1),

I will argue that the correct account of epistemic justification requires at least one psychological capacity for justified beliefs, if not more than one. The conjunction of these claims provides the foundation for the second stage of the argument against (2.1).

The Denial of Skepticism About Justification

The first stage of the argument against (2.1) rests on the rejection of a certain kind of skepticism. The type of argument I will employ for this anti-skeptical thesis has a long and venerable history in epistemology. This approach to skepticism can be found in the writings of Thomas Reid (1970, passim), G.E. Moore (1922, p. 163; 1959, p. 144ff.), and more recently, Gilbert Harman (1973, p. 3-7), Roderick Chisholm (1977, pp. 119-134) and John Pollock (1986, pp. 3-7). The central theme in the attitude of these philosophers toward skepticism is that the considerations advanced on behalf of the thesis of skepticism are always more dubious than the claim that we have knowledge and/or justified beliefs. This treatment of skepticism is often referred to as the "commonsense" tradition in epistemology. Although I will adopt this strategy, my argument requires only the denial of skepticism with respect to the human possession of, and human capacity to possess, justified beliefs (the denial of "J-skepticism(H)"). So my argument rests on a positive answer to question (1.4).

Here is my argument for the denial of J-skepticism(H): "No skeptical argument which concludes with the claim that human beings are incapable of possessing justified beliefs is such that the premises are more obvious than the denial of the conclusion. That is, it is always more reasonable to hold that humans are capable of having justified beliefs than it is to accept any set of premises which entail that humans are not. One reason that it is clear that humans are capable of having justified beliefs is that it is obvious that many humans actually have some justified beliefs, and having justified beliefs entails being capable of having them. Thus, J-skepticism(H) is false."

This rejection of skepticism about human justified beliefs does not involve the dismissal of every form of epistemological skepticism. A more serious form of skepticism applies to knowledge. Whereas the skepticism concerning epistemic justification which is being excluded here amounts to the claim that humans are incapable of having justified beliefs, skeptical arguments about knowledge often rely on the logical possibility of widespread false belief. So the possibility of skepticism about knowledge remains even when skepticism about justification is false. Consequently, the position I adopt here shares with the commonsense tradition the rejection of J-skepticism(H), but not the rejection of skepticism with respect to knowledge. As we shall see, this denial of J-skepticism(H) has interesting

consequences concerning the connection between epistemology and psychology which have gone largely unnoticed.

Many philosophers, such as Goldman, do not share the above approach to skepticism. In (1986, p. 39), Goldman states that whether it is humanly possible to have justified beliefs "hinges on the availability or nonavailability of sufficiently reliable (cognitive) processes." He says this because, on his account of epistemic justification, the possession of sufficiently reliable cognitive processes is a necessary condition for having justified beliefs. What this indicates about Goldman's methodology is that, for him, the question "What is the nature of epistemic justification?" (question (1.3)) has primacy over the question "Do human beings have the capacity for justified beliefs?" (question (1.4)). Given this methodological priority, the question concerning the truth of J-skepticism(H) is left to psychological investigation. If psychological investigation determines that our cognitive processes are sufficiently reliable, then J-skepticism(H) is false. If psychological investigation concludes that these processes are not sufficiently reliable, then J-skepticism(H) is true.

The approach I have taken above differs from Goldman's by reversing the methodological priority of these two questions. For me, the question "Do humans have the capacity for justified beliefs?" must be answered before answering the question concerning the nature of epistemic justifi-

cation. Since my answer to the former question is "yes", I am able to generate a constraint on the answer to the latter question, as we shall see. For Goldman, psychological investigation is required for answering the former question concerning J-skepticism(H), but is not required for answering the latter question concerning the nature of epistemic justification. For me, as we shall see, psychological investigation is required for answering the question about the nature of justification but not required to settle the issue of J-skepticism(H). This difference of methodology explains why Goldman endorses (2.1) and I deny (2.1).

Given these two opposed epistemological methodologies, are there reasons for preferring one of them over the other? There is at least one good reason for choosing the methodology I have followed instead of opting for Goldman's. The following methodological principle seems self-evident to me: "When constructing a theory, one ought to begin with the most obvious relevant considerations." It seems clear to me that it is more obvious that humans have justified beliefs and consequently are capable of having justified beliefs than it is that any particular account of epistemic justification is correct. A non-philosopher would quickly grant that humans have justified beliefs but would be hard-pressed to specify the nature of epistemic justification. Even among philosophers there seems to be much more agreement that humans have justified beliefs than that any particular

account of epistemic justification is correct. There is a notorious lack of agreement in this latter category. I conclude that, given the above methodological principle, my methodological approach is preferable to Goldman's.

The Constraints of Realism and Psychological Realism

Given this denial of J-skepticism(H), we can generate a criterion of adequacy for accounts of epistemic justification. This is the thesis that no account of epistemic justification is correct which entails that ordinary human beings are incapable of having justified beliefs. That is, no account of epistemic justification which entails J-skepticism(H) is true. No such view can be true if J-skepticism(H) is false, and J-skepticism(H) is false. Another way of stating this criterion of adequacy on accounts of epistemic justification is to say that the correct account of epistemic justification must be realistic with respect to human beings (realistic(H)). That is, it must not require of human beings for the possession of justified beliefs any capacity of any kind which human beings do not possess. An account of epistemic justification is false if it ascribes to ideal human beings the capacity for justified beliefs but denies that real human beings have this capability. It is in this somewhat ordinary, non-philosophical sense that I will be using the word 'realistic.' If J-skepticism(H) is false, then humans are capable of having justified beliefs

(humans are "J-capable"). If humans are J-capable, then the correct account of epistemic justification must be realistic(H). Since J-skepticism(H) is false, the correct account of epistemic justification must be realistic(H). Let us call this constraint "the constraint of realism." The constraint of realism can be stated as follows:

(2.2) The correct account of epistemic justification must be realistic(H).

According to the constraint of realism, accounts of epistemic justification must not require for the possession of justified beliefs any capacity of any kind which human beings do not possess. That is, for any account of epistemic justification A and for any type of capacity C, if A requires C for the possession of justified beliefs, then A must be C-realistic(H) (i.e., humans must have C). In this quantified principle, C ranges over an infinite number of types of capacities. Obviously, for the purpose of evaluating plausible candidate accounts of epistemic justification, we will not be interested in every possible type of capacity over which C ranges. Given our interest in answering question (1.6), which concerns the relation between determining the nature of epistemic justification and psychological investigation, we can narrow our focus to the following constraint which is entailed by (2.2). The constraint applies only to psychological capacities:

(2.3) The correct account of epistemic justification

must be psychologically realistic(H).

I will refer to (2.3) as "the constraint of psychological realism." Since (2.3) is entailed by (2.2), it is also supported by the argument for (2.2), given above, which depends on the rejection of J-skepticism(H).

The importance of (2.3) is a function of the likelihood that the correct account of epistemic justification has a psychological component, that is, that it requires some psychological capacity for the possession of justified beliefs. This becomes apparent when it is realized that any account of epistemic justification which has no psychological component satisfies (2.3) by default. If no plausible account of epistemic justification were to incorporate a psychological requirement, then there would never be a need to expend any mental effort to determine whether any plausible account of epistemic justification failed to satisfy the constraint of psychological realism. So the question before us, with respect to (2.3), is this. Given that (2.3) is true, is it philosophically significant? This question reduces to the following question: Does the correct account of epistemic justification have a psychological component? Does one need to possess at least one psychological capacity in order to have justified beliefs?

Epistemic Justification and Psychological Capacities

It is clear that there is at least one psychological

capacity which must be required for justified belief by an account of epistemic justification in order for that account to be acceptable. This is the capacity to have beliefs. One cannot have a justified belief unless one has a belief. Consequently, one cannot be capable of having justified beliefs unless one is capable of having beliefs. So every acceptable account of epistemic justification incorporates at least one psychological requirement for the possession of justified beliefs. Now this fact makes (2.3) more significant than it would be if no plausible account of epistemic justification incorporated any psychological requirement at all for justified belief. Given this required psychological component, no plausible account of epistemic justification will satisfy the constraint of psychological realism by default. However, I take it to be quite obvious that human beings have beliefs and are ipso facto capable of having beliefs. Hence, if the capacity for belief-possession were the only psychological requirement incorporated by plausible accounts of epistemic justification, then although (2.3) would not be satisfied by default, it would be obviously satisfied, and its significance would still be in doubt. Consequently, (2.3) is significant only if some interesting psychological capacity other than the capacity for having beliefs is required for justified belief.

Before we consider the prospects for such an additional psychological capacity, I would like to make a paren-

thetical observation. Since the capacity for having justified beliefs entails the capacity for having beliefs, the argument for (2.2) (and hence for (2.3)) depends for its soundness not only on the rejection of J-skepticism(H), but also on the rejection of eliminativism with respect to beliefs. Recall that my main reason for concluding that humans are capable of having justified beliefs was that it is obvious that we actually have justified beliefs. But if eliminativism were true, then we would not and could not have justified beliefs since we would not have any beliefs. I am inclined to treat arguments for eliminativism in the same manner that I treated arguments for J-skepticism(H). I cannot imagine any grounds which would be so compelling as to convince me that I do not have any beliefs. But even if some convincing argument for eliminativism were forthcoming, this would not necessarily undermine the argument for (2.2) and (2.3). As long as the argument for eliminativism did not exclude the possibility of there being some mental states with content other than beliefs, the argument for (2.2) and (2.3) could be recast in terms of some other contentful mental state besides belief. Although a successful eliminativism excluding all contentful mental states would indeed undermine the argument for (2.2) and (2.3), such a view seems so preposterous as not to be worth considering.

The question with which we were concerned before this brief excursus on eliminativism involved the signifi-

cance of (2.3) as a constraint on accounts of epistemic justification. It is clear that (2.3) is important only if some psychological capacity is required for the possession of justified beliefs in addition to the capacity for belief. Does the correct account of epistemic justification have a psychological component in addition to this minimally necessary one? Recently, many epistemologists have given an affirmative answer to this question. For instance, Alvin Goldman (1979; 1986) has championed a "causal process" view of epistemic justification according to which a belief is justified only if it is produced or sustained by a reliable causal cognitive process. Hence, on Goldman's view, in addition to being capable of having beliefs, one must have causal psychological processes which produce and sustain one's beliefs in order to have justified beliefs. Moreover, these processes must be reliable (this latter necessary condition is more complicated than this in (1986)). I will evaluate this position in chapter five before offering my own proposal for an answer to question (1.3). For now, I am interested in merely pointing out the existence of views which require for justified belief psychological capacities other than the capacity to have beliefs.

In chapter three, we will examine two other views which incorporate a psychological requirement for justified belief. In addition, in chapter four we will look at both a psychological thesis which seems to be presupposed by many

epistemological views, including accounts of epistemic justification, and a thesis about introspection, a psychological process the exercise of which is required by many accounts of epistemic justification for justified belief. In those two chapters we will observe how, among other similar positions, epistemological views which have been offered by Rene' Descartes, W.V.O. Quine, Roderick Chisholm and Keith Lehrer incorporate one or more of these psychological components. Although each of these accounts has a psychological component requiring a certain psychological capacity for the possession of justified beliefs or knowledge in addition to the capacity for belief-possession, it may not be the case that this was intended by each of these philosophers. Not all psychological components are intended and explicitly stated in accounts of epistemic justification. Whether such components are explicit or implicit, intended or unintended, I believe that many influential views of epistemic justification and propositional knowledge in the literature contain them. That prominent philosophers have proposed and defended such views is good prima facie evidence that the correct account of epistemic justification contains an interesting psychological component. At any rate, the existence of such views demonstrates that there is an important role for the constraint of psychological realism to play in the testing of candidate accounts of epistemic justification.

An Objection

One might object to this claim concerning the importance of the constraint of psychological realism by arguing as follows. "No existing account of epistemic justification is initially plausible enough to satisfy a constraint of intuitiveness. If this is the case, then there are sufficient grounds for rejecting all existing accounts of epistemic justification in the absence of applying the constraint of psychological realism. Furthermore, if the constraint is not being employed, there is no important role for it to play in epistemology. Hence, there is no important role for the constraint of psychological realism to play in epistemology".

There are a number of problems with this objection. In particular, although the second premise is true, the first and third premises are arguably false. Consider the first premise. The claim is that no existing account of epistemic justification is plausible enough to satisfy a constraint of intuitiveness. There are at least two reasons for rejecting this assertion. The first is that, when the nature of intuitiveness is properly understood, the premise is obviously false. Intuitiveness is a relative notion. A view is intuitive or unintuitive only from the standpoint of a particular individual. Furthermore, intuitions often vary among individuals and even with respect to the same individ-

ual at different times. What one person takes to be entirely implausible may be and often is considered by another person of equal intelligence and reasonableness quite intuitive. For instance, although many may find the epistemological views which I will consider in the next two chapters unintuitive, it is quite reasonable to suppose that, at least at the time each philosopher proposed and defended his view, he considered it to be intuitive. Furthermore, it is likely that these influential positions have been deemed plausible by a good many other philosophers. Some may continue to endorse them or similar views. Given this relative nature of intuitiveness and the obvious fact that the intuitions of competent philosophers differ over the plausibility of various epistemological positions (with many philosophers finding such views intuitive), it is false to say that no existing account of epistemic justification is intuitive. If the objector has an absolute notion of intuitiveness in mind, one wonders what justification could be offered for making such a claim. Furthermore, it would be clearly illegitimate to argue that no existing account of epistemic justification is intuitive in the absolute sense (if there is one) on the basis of there being none which is intuitive from the standpoint of the critic.

The second reason for denying the first premise of the objection is based on rejecting what is probably the best reason for accepting it. One could argue for this pre-

mise as follows: "If one were able to collect all of the objections in the literature which have been raised against existing accounts of epistemic justification and have the entire philosophical community scrutinize them, the consensus would be that none of these accounts is intuitive. If this were the case, then we would have good grounds for supposing that no existing account of epistemic justification satisfies a constraint of intuitiveness." Now, even if the practical obstacles involved in bringing about the state of affairs described in the premises of this argument could be overcome, the envisioned result would still be extremely unlikely given intuitional relativity and intuitional variability. However, even if we waive these difficulties, there is a further reason for rejecting this argument for the first premise. This is the observation that, as will become apparent in the next two chapters, some, if not many epistemological views which are rejected in the name of implausibility are really discarded on the basis of psychological unreality and not on the basis of strict unintuitiveness. If this is the case, then the so-called implausibility of an account of epistemic justification would not always be a reason for thinking that the account is unintuitive in a strict, narrow sense. Neither, therefore, would it be a reason for thinking that the constraint of psychological realism had no important role to play in its evaluation. A conclusion of implausibility may be

reached on the basis of applying the constraint itself in such a case.

I believe that the previous two paragraphs provide sufficient grounds for rejecting the first premise in the argument for the insignificance of the constraint of psychological realism. Consequently, there is no good reason to accept the conclusion of that argument. Moreover, even if the first premise were granted, the conclusion could still be avoided. This is because the third premise is also false. According to the third premise, if the constraint of psychological realism is not currently being employed, then there is no important role for it to play in epistemology. The following considerations demonstrate the falsity of this claim.

Even if no existing account of epistemic justification were to satisfy a constraint of intuitiveness, so that the constraint of psychological realism would be superfluous relative to all existing accounts, it would still be important relative to any as-yet-undiscovered accounts which would satisfy some constraint of intuitiveness. Hence, its actual failure to be used would not entail its lack of importance with respect to all future epistemological concerns. Furthermore, if the argument for the constraint of psychological realism is correct, the correct account of epistemic justification, if there is one, will satisfy the constraint. It is reasonable to suppose that the discovery

of the correct account of epistemic justification could not be confirmed without applying the constraint of psychological realism. So this third premise in the argument for the insignificance of the constraint of psychological realism is false. Consequently, the case against this objection is overdetermined. I conclude that, in the absence of another, compelling objection, there is no good reason to doubt the significance of the constraint of psychological realism.

Summary

This completes the first stage of the argument against (2.1), the claim that it is possible to determine the correct account of epistemic justification apart from psychological investigation. In this stage of the argument, I have defended the view that the correct account of epistemic justification must be psychologically realistic(H). That is, it must not require for the possession of justified beliefs some psychological capacity which human beings do not possess. The argument for this conclusion rests on the rejection of J-skepticism(H) and eliminativism with respect to beliefs. I have also provided grounds for thinking that the methodological strategy which underlies my endorsement of this argument for the constraint of psychological realism is preferable to an opposing strategy. Moreover, I have offered reasons for thinking that this constraint has an important role to play in the assessment of candidate

accounts of epistemic justification. The argument for this latter thesis is based on the observation that, not only do all plausible views of epistemic justification require the capacity for belief-possession for justified beliefs, it is also the case that many influential views require additional psychological capacities for justified belief as well. Finally, I have argued that an objection to this latter claim concerning the importance of the constraint of psychological realism does not succeed.

Stage Two: The Need for Psychological Investigation

The above conclusions prepare the way for the second and final stage of the argument against (2.1). The conclusion of this second stage is that it is not possible to determine whether an account of epistemic justification is psychologically realistic apart from psychological investigation. Given the main conclusion of the first stage, that the correct account of epistemic justification must be psychologically realistic, it seems quite reasonable to suppose that it would not be possible to determine the correct account of epistemic justification without determining whether the account is psychologically realistic. It follows that it is not possible to determine the nature of epistemic justification (i.e., give an account of epistemic justification) apart from psychological investigation. That is, on the basis of these considerations, we can conclude

that (2.1) is false. However, we have yet to see the argument for the claim that testing an account of epistemic justification for psychological realism requires psychological investigation.

The Indispensability of Psychology

Here is the argument for this crucial claim in the case against (2.1). I will refer to this argument as "the indispensability argument" and to its conclusion as "the indispensability thesis." According to the constraint of psychological realism, an account of epistemic justification must not require for the possession of justified beliefs a psychological capacity which humans do not have. In order to determine whether a candidate account of epistemic justification satisfies this constraint, one must know what cognitive capacities humans possess. This kind of information can be acquired only by means of psychological investigation. Psychological investigation consists in either the methods of experimental cognitive psychology or such non-scientific methods as introspection. Neither of these methods alone is sufficient to answer all the relevant psychological questions. Consequently, both introspection and experimental cognitive psychology are indispensable tools in the determination of the correct account of epistemic justification.

I will not discuss the claim concerning the insuffi-

ciency of either experimental psychology or introspection alone until chapter three. For now I am interested only in the claim that psychological investigation, of some kind or other, is indispensable in testing an account of epistemic justification for psychological reality. What objections can be raised against this view? I will consider three counterarguments, in the order of their importance. None of these arguments, as far as I can see, undermines the conclusion of the indispensability argument. In responding to the third objection, I introduce a distinction which not only contributes to showing why the objection fails, but also helps to explain the existence in the literature of psychologically unrealistic accounts of epistemic justification, some of which will be examined in chapters three and four.

Three Objections to the Indispensability Argument

Objection #1. One might object to the thesis of the indispensability of psychological investigation by pointing out, as I mentioned earlier, that accounts of epistemic justification with no psychological component whatsoever automatically satisfy the constraint of psychological realism. In cases such as these one can tell immediately that the account satisfies the constraint without doing any psychological investigation. Hence, psychological investigation is not always required for determining whether an account of epistemic justification is psychologically realistic.

Although this is indeed the case, the point can be granted without the need to make any significant concessions. As I argued in the previous section, all minimally plausible accounts of epistemic justification must have at least one psychological component. One cannot have justified beliefs without the capacity for possessing beliefs. Furthermore, we saw that there are good prima facie grounds for concluding that there is at least one additional more interesting psychological requirement for justified belief. Accordingly, we can exclude accounts of epistemic justification which do not have a psychological component at the outset without testing them for psychological reality. So, practically speaking, the question concerning the psychological reality of such accounts does not arise. In light of these considerations, the thesis concerning the indispensability of psychological investigation can be revised to obviate the objection: Psychological investigation is required in order to test accounts of epistemic justification which have at least a prima facie minimal plausibility. For convenience, this qualification will be suppressed in the discussion to follow.

Objection #2. Another objection to the thesis of the indispensability of psychological investigation in testing accounts of epistemic justification for psychological reality focuses on accounts which have as their sole psychological requirement for justified belief the capacity for

having beliefs. Suppose that we are wondering whether such an account satisfies the constraint of psychological realism. The objector might argue as follows: "It is obvious that human beings have beliefs. Hence, it is equally obvious that humans are capable of having beliefs. If it is obvious that humans are capable of having beliefs, then psychological investigation is not needed to confirm that they are. Therefore, psychological investigation is not required for determining whether humans are capable of having beliefs. If this is the case, then psychological investigation is not always required for testing even initially minimally plausible accounts of epistemic justification for psychological reality."

This argument fails in virtue of the falsity of the claim that psychological investigation is not needed to confirm the obvious thesis that humans are capable of having beliefs. Regardless of how obvious this thesis is, it is clearly the case that we know it on the basis of introspection and the informal observation of other human beings. Although neither of these methods is scientific in the narrow sense, each is a legitimate method of psychological investigation in the rather broad sense in which I am using this term. Furthermore, it is difficult to imagine a basis for knowing that humans are capable of belief-possession, which is available to us, that does not ultimately incorporate at least a partial appeal to psychological investiga-

tion. Consequently, this objection to the thesis of the indispensability of psychological investigation does not succeed.

Objection #3. The final and most important objection was suggested to me by John Pollock (in conversation). If the argument for the indispensability of psychological investigation is sound, then knowing that an account of epistemic justification is intuitive is not sufficient for knowing that it is psychologically realistic. This is because semantic intuitions can yield only a priori knowledge, whereas, if the indispensability argument is sound, a posteriori knowledge is required for testing an account of epistemic justification for psychological reality. Pollock's argument purports to show that intuitiveness is always sufficient for knowing that a view of justification is psychologically realistic. The argument proceeds as follows. "If an account of epistemic justification is intuitive, then there is good reason to think that it is true. If there is good reason to think that it is true, then there is good reason to think that it is consistent with all psychological truths, since, of course, all truths are consistent. If there is good reason to think that an account of justification is consistent with all psychological truths, then there is good reason to believe that it is psychologically realistic. Consequently, if an account of epistemic justification is intuitive, then there is good reason to think

that it is psychologically realistic."

In response to this objection, I would like first to point out that there are at least two ways in which the word 'intuitive' can be understood. I have hinted at this distinction already in this chapter. When I considered the argument to the effect that the constraint of psychological realism is unimportant given the false claim that no existing accounts of epistemic justification are intuitive, I suggested that what is taken to be an application of a constraint of intuitiveness is sometimes, if not often, an application of the constraint of psychological realism. I believe that what explains this is that 'intuitiveness' is being used in these cases in a rather broad sense to include both semantic intuitions and the deliverances of introspection. A narrower sense of 'intuitiveness' includes only appeal to semantic intuitions. When I talk about a constraint of intuitiveness, I intend the word 'intuitiveness' in this latter narrow sense. Hence, when a judgment of "unintuitiveness" is reached about an account of epistemic justification on the basis of introspection, the broader sense of the word is operative and the judgment is based on an application of the constraint of psychological realism.

Is Pollock's objection best construed as appealing to the narrow sense of 'intuitiveness' or the broad sense? If 'intuitiveness' is being used here in the broad sense, and the appeal to introspection is doing all the work, then

although the argument interpreted in this way is quite plausible, it is not an objection to the indispensability argument. This is because it would not show that it is possible to determine that an account of epistemic justification is psychologically unrealistic apart from psychological investigation, since introspection is itself a form of psychological investigation. This understanding of Pollock's argument would, however, be a challenge to my contention that introspection is not always sufficient as a test for psychological reality. As I previously said, I will take up this issue in chapter three.

If Pollock's objection is to succeed against the indispensability thesis, 'intuitiveness' must be taken in the narrow sense. Before we examine the prospects for his argument construed in this way, let us examine a thesis which is similar to but stronger than his conclusion. A consideration of this similar thesis will put us in a better position to evaluate Pollock's argument interpreted in terms of the narrow sense of 'intuitiveness'. The conclusion of Pollock's argument is that the intuitiveness of an account of epistemic justification would provide us with good reason for thinking that the account is psychologically realistic. The similar but stronger claim is that all intuitive accounts of epistemic justification are psychologically realistic. This latter claim is easily seen to be false. It is surely at least logically possible that an account of

epistemic justification be considered intuitive which is nonetheless psychologically unrealistic. That is, it is conceivable that there be an account of epistemic justification which explains and unifies intuitive judgments about justified belief while at the same time requiring for the possession of justified beliefs some psychological capacity which humans do not possess. But this might seem rather odd at first glance. A psychologically unrealistic account of epistemic justification entails J-skepticism(H). Someone who found such an account intuitive would be in the interesting position of finding an account intuitive which entailed the quite unintuitive result of J-skepticism(H). But how could an intuitive account have such an unintuitive result? Furthermore, if a person who denies J-skepticism(H) accepts a psychologically unrealistic account, then the person accepts a proposition which entails the denial of another proposition the person accepts. Does this count against the conceivability of my claim that an intuitive account can be a psychologically unrealistic account?

Some of the apparent oddness of a person finding an account intuitive which leads to a result the person would find unintuitive is diminished when it is pointed out that at least two kinds of intuitiveness are operative in such situations. The intuitiveness a person attributes to an account of epistemic justification includes semantic intuitions. On the other hand, the intuitiveness which someone

might attach to the denial of J-skepticism(H) has nothing to do with semantic intuitions. It is based instead on a conviction about the existence of justified beliefs. Given that there are two sources of conviction in such cases, instead of one, it seems less odd that a person could endorse a proposition which entails the denial of another proposition the person endorses. This is because the person would be endorsing these propositions on entirely separate grounds. However, this observation about distinct sources of intuition is not completely adequate as an explanation of the phenomenon the possibility of which is in doubt. What other explanation can be offered?

I believe that these cases of intuitive account entailing an unintuitive result can be explained in such a way as to show that they are not only conceivable, but likely to happen. As I mentioned above, accounts of epistemic justification can incorporate psychological components either implicitly or explicitly. A psychological component is implicit when it is not explicitly stated in the account but is entailed by what is explicitly stated in the account. For instance, that the capacity to have beliefs is required for the possession of justified beliefs is often implicit in an account of epistemic justification. Although it is difficult to imagine a person failing to notice this implicit psychological component when evaluating an account of epistemic justification, it is easy to imagine someone overlook-

ing a psychological component which is entailed less obviously by an account. We will see an example of this in the next chapter. If it is possible for someone to evaluate an account of epistemic justification without seeing all of its implicit components, then it seems possible for someone to find an account intuitive which is also psychologically unrealistic. Thus, if a psychological component of an account of justification is implicit and not obvious, then a person could find the account intuitive even though it leads to a skeptical conclusion which the person would deem unintuitive. What about accounts of epistemic justification which have explicit psychological components? The preceding explanation fails to explain how someone could find this kind of account intuitive when it is also psychologically unrealistic. This kind of case requires a different explanation.

The explanation which explains this possible response to accounts with explicit psychological components employs a distinction between "conditional intuitiveness" and "unconditional intuitiveness." A proposition is conditionally intuitive if and only if its intuitiveness is at least partially a function of a conscious or unconscious psychological assumption. A proposition is unconditionally intuitive if and only if it is not conditionally intuitive. This distinction can be used to explain how a person could think that an account of epistemic justification with an

explicit psychological component is intuitive even though it entails J-skepticism(H) (a thesis the person rejects). If a person assumes, consciously or unconsciously, that humans have a certain psychological capacity which they nevertheless do not possess, then an account of justification which requires this capacity may very well seem intuitive to that person. This is because it may explain and unify the person's intuitive judgments about possible cases of justified belief. The intuitiveness of these judgments and of the overall account may depend partially on the psychological assumption. At least, the person's lack of recognition that the assumption is false may prevent her from seeing the skeptical consequences of her view. For instance, if one believes that humans are capable of examining all the contents of their minds at once, then it may not seem unintuitive to one that such a cognitive operation is required for justified belief. Of course, humans are not capable of this, and if such an ability were required for the possession of justified beliefs, no human being would have any justified beliefs. What explains how intuition could endorse an account with a false explicit psychological component such as this is that the account is only conditionally intuitive. Discovery of the psychological unreality of the account would undermine this intuitiveness, since it would eliminate the psychological assumption upon which the intuitiveness depends.

I conclude that, whether an account of epistemic justification has an implicit psychological component or an explicit psychological component, it is at least logically and most likely at least psychologically possible for an intelligent and reasonable person who finds J-skepticism(H) unintuitive to find such an account intuitive even though it is psychologically unrealistic. Consequently, the claim that all intuitive accounts of epistemic justification are also psychologically realistic is false.

Given the falsity of this thesis, what are the prospects for Pollock's argument which concludes with the weaker claim that if an account of epistemic justification is intuitive, then there is good reason to believe that it is psychologically realistic? Recall that if Pollock's argument is to succeed against the indispensability thesis, 'intuitiveness' must be understood in the narrow sense, that is, in the sense which involves appeal only to semantic intuitions. There is a further ambiguity in the argument which needs to be made explicit in order to evaluate it properly. What is the force of 'good reason'? This term can mean either 'sufficient reason' or 'some (less than sufficient) reason'. Let us consider the merits of Pollock's argument on both interpretations.

Let us first examine the version of the argument one gets by construing 'good reason' as 'sufficient reason'. If this version of the argument is sound, it surely defeats

the indispensability thesis. If intuition provides us with sufficient grounds for determining whether an account of epistemic justification is psychologically realistic, then resorting to psychological investigation for this purpose is unnecessary. The problem with this version of Pollock's argument is that its first premise is clearly false. The first premise of this version is that if an account of epistemic justification is intuitive, then there is sufficient reason to think that the account is true. The falsity of this claim becomes apparent when the various features of intuitiveness (in the narrow sense), which we have already discussed, are kept in mind. These properties of intuitiveness make it less than sufficient as an indication of the truth of a proposition.

In the first place, there is the thesis of intuitional relativity. This is the claim that all statements of the form 'P is intuitive' (where 'P' is a schematic letter to be replaced by constants referring to propositions) are elliptical. All such statements have an implicit index to an individual person and a particular time. The full reading of statements of this form is: 'For some person S and time t, P is intuitive to S at t'. Intuitional relativity leaves open the possibility of two people having conflicting intuitions about the truth of a certain proposition at the same time and one person having one intuition about a given proposition at one time and a conflicting intuition about it

at a different time. The realization of either of these possibilities constitutes intuitional variability. Moreover, it is clear that intuitional variability is often the case among philosophers. A survey of current epistemological literature would confirm this. Given this, although the intuitiveness of an account of epistemic justification may be some reason to think that the account is true, it cannot be sufficient reason. This can be made clear by the following example.

Suppose that account A of epistemic justification is inconsistent with account B. Further, suppose that person S finds A intuitive and person T finds B intuitive. If we assume that intuitiveness provides a sufficient reason for thinking that an account of epistemic justification is true, then it would follow that there is sufficient reason for thinking that A is true and sufficient reason for thinking that B is true. But this is absurd, since A is inconsistent with B. It is more reasonable to say in such cases that intuitiveness provides some reason for thinking that an account of epistemic justification is true. So S has some reason to believe that A is true and T has some reason to think that B is true.

If situations such as the one in the previous paragraph were merely possible and not actual or rarely actual, then it would not undermine the contention that intuitiveness in the narrow sense is always a sufficient reason for

believing that an account of epistemic justification is true. However, it is clear that such cases abound in the philosophical literature. For instance, Goldman (1986, p. 116) maintains that epistemic justification has an interesting connection with truth. On the other hand, Pollock (1986, p. 183) insists that ordinary epistemic justification does not have an interesting connection with truth. I presume that both of these philosophers find their accounts intuitive. Given this, it would be unreasonable to conclude that the intuitiveness of an account of epistemic justification provides a sufficient reason for thinking that it is true.

I conclude that the first premise of Pollock's argument against the indispensability thesis is false when that argument is construed as being about intuitiveness in the narrow sense and when 'good reason' is taken to mean 'sufficient reason'. Hence, that form of the argument, which would undermine the indispensability argument if it were sound, is unsound. Would Pollock's argument succeed if 'good reason' is interpreted as 'some (less than sufficient) reason'? I will grant that this form of the argument is sound. However, its conclusion falls short of defeating the indispensability argument. This is because there is more than just the mere psychological possibility that a conditionally intuitive account is psychologically unrealistic. The next two chapters will illustrate how intuitions can

lead even those who would deny J-skepticism(H) to adopt a psychologically unrealistic account of epistemic justification. After all, what reason is there to think that substantive empirical questions can be answered correctly on the basis of relatively unreliable semantic intuitions? I conclude that Pollock's objection fails on any reasonable interpretation of its premises and conclusion.

Summary

In this section I have argued that psychological investigation, broadly construed, is an indispensable tool in determining whether an account of epistemic justification is psychologically realistic. I have considered three objections to this thesis. The first, least important objection was that since accounts of epistemic justification without psychological components satisfy the constraint of psychological realism by default, psychological investigation is not always required for determining whether an account of epistemic justification is psychologically realistic. This led to a qualified version of the indispensability thesis: All initially, minimally plausible accounts of epistemic justification must be tested for psychological realism by means of psychological investigation. The second objection turned on the obviousness of the thesis that the capacity to have beliefs is required for the possession of justified beliefs. This obviousness was held to eliminate

the need for psychological investigation when testing accounts with only this psychological component for psychological reality. But, obvious or not, such a thesis can be verified only on the basis of psychological investigation. Finally, the most important objection to the indispensability thesis, due to Pollock, held that the intuitiveness of an account of epistemic justification would be good reason to conclude that the account was psychologically realistic. I argued that, on one possible construal of this argument, although its soundness would entail the denial of the indispensability thesis, there is good reason to think that it is unsound. Furthermore, on the only other reasonable reading of this argument, although it is sound, its conclusion does not undermine the indispensability thesis. Given the failure of these objections and the plausibility of the indispensability argument, I conclude that the indispensability thesis is true.

Conclusion

As I said above, if the indispensability thesis is true, then given the constraint of psychological realism and the assumption that one cannot determine the correct account of epistemic justification unless one determines that it is psychologically realistic, it follows that (2.1) is false. That is, it follows that one cannot determine the correct account of epistemic justification apart from psychological

investigation, broadly construed. This argument bears repeating in another form: "The correct account of epistemic justification must be psychologically realistic. If the correct account of epistemic justification must be psychologically realistic, then one can determine the correct account of epistemic justification only if one determines that the account is psychologically realistic. Hence, one can determine what the correct account of epistemic justification is only if one determines that the account is psychologically realistic. One can determine that a candidate account of epistemic justification is psychologically realistic only if one makes use of psychological investigation, broadly construed. Hence, one can determine what the correct account of epistemic justification is only if one makes use of psychological investigation." This conclusion is the thesis of epistemological naturalism.

What are the implications of this thesis for the argument given on behalf of (2.1) in the introduction to this chapter? Obviously, if the argument for (2.1) is valid, and I assume that it is, one or more of the premises must be false. Recall that the argument for (2.1) has three premises. The first premise equates giving an account of epistemic justification with analyzing the concept of epistemic justification. The second premise holds that conceptual analysis is an entirely a priori enterprise. Finally, the third premise is that if an enterprise is entirely a

priori, that is, if it needs to make use of only a priori considerations, then it will not need to make use of any a posteriori considerations. Now I take it that the third premise is necessarily true. This means that, if the thesis I have defended in this chapter is correct, it is either the case that giving an account of epistemic justification is not just to analyze the concept of epistemic justification, or it is the case that conceptual analysis is an activity that is not entirely a priori. Of course, both of these alternatives might be true. Although it would be interesting to pursue this question, I will leave that project for another occasion. It is enough to point out at this juncture that the main thesis of this chapter has some quite interesting philosophical consequences.

CHAPTER THREE

APPLYING THE CONSTRAINT OF PSYCHOLOGICAL REALISM:

NON-SCIENTIFIC PSYCHOLOGY

In chapter two, I argued that one cannot determine the correct account of epistemic justification apart from psychological investigation. My reasons for this claim are that the correct account of epistemic justification must be psychologically realistic and that one cannot test an account of epistemic justification for psychological reality apart from psychological investigation. The main purpose of this chapter and the next is to illustrate the application of the constraint of psychological realism. In order to do this, one need only consult current epistemological literature. Some epistemologists have employed the constraint of psychological realism in their critical evaluations of the epistemological views of other philosophers. In the next two chapters, I will survey two recent applications of variations of the constraint as defined in chapter two (that is, as applying to accounts of justification), one application of a version of the constraint applied to theories of rationality, and an argument which has a structure similar to arguments employing the constraint and which is suggestive of further applications of it. Although these latter two examples are variations on the theme of employments of the

constraint of psychological realism, for simplicity I will refer to all four illustrations as applications of the constraint of psychological realism. However, I will point out the differences among them as we go along.

The four examples which we will examine are the following. First, we will take a look at William Alston's (1985, pp. 64-65) criticism of some remarks of Roderick Chisholm (1977, p. 14) about epistemic justification. Second, we will review Stewart Cohen's (1984, pp. 285-288) critical evaluation of Keith Lehrer's analysis of epistemic justification in Knowledge (1974, p. 198). Next, the discussion will focus on Christopher Cherniak's (1986, pp. 49-71) scrutiny of a feature of Willard van Orman Quine's program of naturalizing epistemology (1960, p. 11; 1961, pp. 42-44). Finally, we will turn to Hilary Kornblith's arguments against an aspect of traditional and contemporary Cartesian epistemology in a recent unpublished work.

Overview of the Illustrations

Before we proceed with a detailed examination of these four examples of employing the constraint of psychological realism, it will be helpful to have before us a general overview of the illustrations. This overview will have two main parts. First, we will look at three ways in which the examples can be categorized in terms of what they illustrate about applying the constraint. These three ways are:

(1) what the examples illustrate about some of the types of accounts of epistemic justification which can be tested by means of the constraint of psychological realism; (2) what the examples illustrate about the types of psychological components incorporated by accounts of epistemic justification to which the constraint can be applied; and (3) what the examples illustrate about the types of psychological investigation to which one can appeal in applying the constraint. This last section will include a rather extended discussion of the respective roles of non-scientific and scientific psychological investigation in applying the constraint of psychological realism. The second main part of the overview will involve a brief description of the general structure shared by each of the four critical arguments.

After the overview, the rest of this chapter will be taken up with Alston's treatment of Chisholm and Cohen's discussion of Lehrer. In chapter four, we will continue our illustration of ways in which the constraint of psychological realism has been and can be applied with Cherniak's arguments against Quine and Kornblith's scrutiny of Descartes and contemporary Cartesian epistemology. Chapter four will end with an observation of three general consequences which would follow upon the success of these four applications of the constraint of psychological realism, and others like them. It is likely that, if the four arguments we will consider are sound, then, given the repre-

sentative nature of the epistemological positions which would be undermined, quite a few other similar views would be threatened as well. I will argue that the consequences of such widespread psychological unreality underscore the importance of the constraint of psychological realism and the indispensability of psychological investigation in its application. I will also point out how such a result would provide motivation for a proposal I will make in chapter five.

What the Examples Illustrate about the Constraint

In this section we will look at types of accounts of epistemic justification, types of psychological components and types of psychological investigation.

Types of Accounts of Epistemic Justification. The examples to follow are criticisms of epistemological theses which fall into two general categories. Although these two categories do not exhaust logical space, they do include many current and influential epistemological positions. As we shall see, each of the four critics claims that the view he is attacking is representative of many similar views in the literature. The first category can be described as "Regulative Deontic Internalist Theories of Epistemic Justification" (Another way of describing the views in this category is to say that they analyze epistemic justification in terms of "epistemic responsibility." For convenience, I

will use this locution to refer to theses in this category.). The views in this category are regulative because they make the epistemic justification of a belief a function of the believer's adherence to action-guiding norms. They are deontic because they explicate epistemic justification in deontic terms such as 'is obliged,' 'is permitted' and 'is forbidden.' Finally, they are internalist because they make justification a function of cognitive states which are within the perspective of the believer. Of the four illustrations which we will examine, two focus on epistemological theses which are either within this category or which share most of the features of accounts in this category. One of these is Alston's argument against Chisholm. Alston construes Chisholm's account of epistemic justification as a version of the epistemic responsibility approach to epistemic justification. The other is Kornblith's criticism of a Cartesian epistemological thesis. Kornblith considers the Cartesian tradition to be primarily concerned with epistemic responsibility in the pursuit of truth. Although Kornblith does not examine an account of epistemic justification, the thesis he considers is regulative, deontic and internalist, as we shall see. Kornblith's arguments do, however, have interesting consequences for accounts of epistemic justification in this category.

The second category of epistemological positions to which the constraint of psychological realism has been

applied, as the examples illustrate, can be called "Positive Holistic Internalist Coherence Theories of Epistemic Justification" (for convenience, "coherence theories"). These theories are coherence theories since they make the justificational status of a belief a function of its coherence with other beliefs in a system of beliefs. They are positive because they require for the justification of a belief that the believer be in possession (here and now) of all of the reasons which support the belief. They are holistic because they make the relation 'x is a reason for y' such that, where the value of 'y' is a belief in a system of beliefs, the value of 'x' is all the beliefs in the system (see Pollock 1979, pp. 101-102, for a discussion of positive holistic coherence theories). Given these three properties, we can see that a belief is justified on this view only if there is a relation of coherence among all of one's beliefs which is such that, for each justified belief in the system, the entire system constitutes a positive reason for that belief. These theories are internalist for the same reason that epistemic responsibility views are internalist. The remaining two illustrations are aimed at these types of theories. These two illustrations are Cohen's consideration of Lehrer's coherence theory of epistemic justification and Cherniak's discussion of Quine's epistemological holism.

In addition to their not exhausting logical space, these two categories of accounts of epistemic justification

are not mutually exclusive. Although there are theories which fall into only one or the other category, some theories belong in both categories. For instance, Bonjour's coherence theory of epistemic justification (1985) is a theory which is in both categories. On the other hand, Harman's (1973) account is a coherence theory but is most likely not an epistemic responsibility theory. Moreover, Chisholm (1977) has an epistemic responsibility view but it is not a coherence account. We will investigate some positions in chapter five which do not fall into either category.

Insofar as the four epistemological theses criticized by the aforementioned evaluators are representative of other epistemological positions in these two categories, the success of the evaluators' arguments would extend to these other theories as well. We will consider the consequences of such widespread psychological unreality at the end of chapter four. A theory such as Bonjour's, which is of both types, may be found to be (at least) doubly psychologically unrealistic. In the next section on types of psychological components, we will see that a view such as Bonjour's may contain more than one psychological component which fails to stand up to psychological scrutiny.

Types of Psychological Components. As I mentioned in the previous chapter, each of the epistemological views evaluated by these critics contains an interesting psycho-

logical component over and above a component requiring the capacity for beliefs for the possession of justified beliefs. Here is a brief summary of the psychological component of each account which will be highlighted in this chapter. (1) A natural interpretation of Chisholm's explanation of epistemic justification entails that, in order to have justified beliefs, one needs to exercise a capacity of direct voluntary control over one's beliefs. As we shall see, this is an example of an unobvious, implicit psychological component. This type of psychological component was discussed in chapter two in the context of explaining how a philosopher who rejects J-skepticism(H) could find intuitive a view which is actually psychologically unrealistic. (2) Lehrer's analysis of epistemic justification in Knowledge requires that one exercise a capacity to form certain meta-beliefs about the probability of one's beliefs in order for these latter beliefs to be justified. This psychological component is explicit in Lehrer's account. If his account is psychologically unrealistic, one could explain a person's finding Lehrer's position intuitive by appealing to the notion of conditional intuitiveness which was introduced in chapter three. That is, Lehrer's view could be intuitive to a person who rejects J-skepticism(H) if the person believes, consciously or unconsciously, that humans have the capacity to form the requisite meta-beliefs. (3) Some of Quine's writings suggest that his program for naturalizing episte-

mology presupposes a model of human memory which is psychologically unrealistic. The model requires of humans certain psychological capacities which our best cognitive theories do not ascribe to them. Cherniak believes that this highly idealized Quinean model of memory is "tacitly presupposed in most philosophy" (1986, p. 70). If it is, this particular application of the constraint of psychological realism has far-reaching consequences. (4) Kornblith believes that traditional and contemporary Cartesian epistemology features introspection as playing a crucial role (Unpublished, p. 3). However, he argues that such an epistemological approach must be revised to do without introspection, since, although humans certainly possess this psychological capacity, introspection is not sufficiently reliable for the goals it is called upon to achieve by this approach (p. 4). Kornblith also believes that this Cartesianism "informs much of contemporary epistemology" (p. 4). Again, if this is the case, this employment of an argument similar in structure to arguments incorporating the constraint of psychological realism has wide-ranging implications.

These four psychological capacities each fall into one of three categories which may reasonably be thought to exhaust the kinds of psychological components incorporated by epistemological theses to which the constraint of psychological realism can be applied. These three categories are: (1) belief-states, (2) belief-processes and (3) the belief-

system. Belief-states are just beliefs of various kinds. Cohen's application of the constraint of psychological realism to Lehrer's analysis of epistemic justification is an example of bringing psychological evidence to bear upon the question of the existence of a certain kind of belief-state. In this case the relevant belief-states are the meta-beliefs described above. Belief-processes are mental processes which eventuate in the formation, maintenance or elimination of a belief-state. Alston's application of the constraint of psychological realism to Chisholm's view of epistemic justification illustrates psychological investigation of a thesis about belief-processes. In this case, the thesis concerns the relationship between the will and belief-processes. In particular, are belief-processes subject to voluntary control? Kornblith's discussion also illustrates the use of psychological investigation to test an epistemological thesis which requires for its truth that a psychological thesis about a certain kind of belief-process be true. The belief-process in question here is the formation of beliefs by means of introspection. The question Kornblith tackles concerns the reliability of this belief-process. Finally, the belief-system encompasses all of a person's belief-states and some or all of the person's belief-processes. Cherniak applies a version of the constraint of psychological realism to a model of the human belief-system which appears to be endorsed by Quine. So

Cherniak's discussion illustrates the use of the constraint in testing epistemological theses with psychological components having to do with the human belief-system. In sum, our four examples illustrate applications of the constraint of psychological realism to theses incorporating each of the three kinds of psychological components listed above.

We have isolated only one psychological component of each of the epistemological accounts mentioned above for our subsequent discussion. Moreover, each of these components is entirely distinct from the others. In spite of this, it is clear that each of these views has more than one psychological component and that the views share some of the same psychological components. For instance, it is often held (see for instance Kornblith, unpublished) that Descartes's epistemological views presuppose the capacity for direct voluntary control of one's beliefs. In addition, something like Quine's idealized model of the human belief system may underlie Lehrer's account of epistemic justification in Knowledge. Although we will not be investigating Bonjour's (1985) account extensively below, it deserves brief mention here as an example of a theory of epistemic justification which contains more than one of the psychological components which will be under consideration in this chapter and the next (see 1985, pp. 153-156 for an overview of his theory which reveals the psychological components listed below). Although Bonjour has been accused of endorsing the thesis of

direct voluntary control of belief (see below in the section entitled "Alston on Chisholm") he denies this in (1985, p. 230). However, it is clear that his view incorporates the same type of psychological component concerning meta-beliefs as does Lehrer's view. Furthermore, it is arguable that it requires for the possession of justified beliefs an idealized model of the human belief-system such as the one Chermiak ascribes to Quine. Finally, although it does not require reliable introspection for the possession of justified beliefs, Bonjour's view does rely heavily on introspection. Consequently, if introspection is unreliable for some purposes, as Kornblith argues, although it would not follow from this alone that Bonjour's view leads to skepticism, what might follow is that his account of epistemic justification makes justification quite weak (it may not be very truth-conducive). At any rate, the success of the following four applications of the constraint of psychological realism in undermining their target accounts would have serious consequences for a theory such as Bonjour's which incorporates at least two of the psychological components under consideration.

If the psychological requirements of some of the views examined in this chapter and the next overlap, then they may stand or fall together. In addition, given that each position has more than one psychological component, even if one application of the constraint of psychological

realism should fail to undermine an account, it may be defeated by means of a different employment of the constraint. However, in the discussion which follows, we will restrict our attention to the one psychological component of each of the four views which was mentioned above.

Types of Psychological Investigation. The four example applications of the constraint of psychological realism can be divided into two groups of two, based on the primary method of psychological investigation which is used. On the one hand, Alston and Cohen employ non-scientific methodology, in particular, introspection outside of the context of a controlled experiment. On the other hand, Cherniak and Kornblith appeal to the experimental results of current scientific cognitive psychology. This taxonomy of kinds of psychological investigation can be made more complex by observing that there are roughly two kinds of psychological considerations to which one can appeal in applying the constraint of psychological realism, whether one is engaging in scientific or non-scientific investigation. First, there are experimental considerations. Secondly, there are theoretical considerations.

Experimental considerations are the relatively low-level results of observation. Theoretical considerations are the relatively high-level theses which are entailed by relatively well-developed and comprehensive psychological theories. Thus, with respect to scientific psychological

investigation, one can appeal to either experiments in cognitive psychology and their results or to full-blown scientific psychological theories. Kornblith's use of psychological investigation consists almost entirely in the appeal to experiments in cognitive psychology. On the other hand, Cherniak bases the psychological premises of his argument mostly on current scientific psychological theories. With respect to non-scientific psychological investigation, the analogy to laboratory experiment is both introspection and the informal observation of other persons. In addition, it is widely believed that, corresponding to explicit scientific psychological theories, there are rudimentary, mostly implicit non-scientific or commonsense psychological theories which underlie the ability of humans to explain and predict their own and others' behavior. Both Alston and Cohen employ introspection in their application of the constraint of psychological realism. On the other hand, Cherniak makes use of some theses of commonsense psychological theories in addition to his appeal to scientific theories. In sum, there are both scientific and non-scientific methods of psychological investigation. Both of these methods can involve appeal to either experiment or theory. This appears to exhaust the possible kinds of psychological investigation which can be used in applying the constraint of psychological realism.

We now have before us a taxonomy of kinds of psycho-

logical investigation. However, there is an apparent problem with one kind of method in this taxonomy which needs to be cleared up before we can proceed with our survey of the illustrations. Given the controversial status of introspection as a method of psychological investigation, should we not restrict the application of the constraint of psychological realism to the use of scientific psychological investigation? At this point we cannot and need not impose this limitation on the constraint. This will become clear in light of the following considerations.

Someone might take the position that the constraint of psychological realism ought to be applied only when relevant scientific psychological investigation is available on the basis of holding that non-scientific methods such as introspection are entirely inadequate as methods of psychological investigation. However, this view is too pessimistic about the merits of introspection. I think it is clear that introspection is not entirely worthless as a form of psychological investigation. Moreover, it seems reasonable to suppose that, for some purposes, it is quite adequate as a psychological method. If this is granted, it could still be argued that, since scientific methods are always preferable, the employment of the constraint of psychological realism ought to be restricted to scientific psychological investigation. There are two reasons to reject this position. One is that, even if scientific methods are always

preferable, they are not always available. So we cannot restrict the use of the constraint of psychological realism to scientific methods alone. The second is that scientific methods are not always preferable to non-scientific methods. Consequently, we need not restrict the constraint to scientific methods.

Here are some considerations which support the first reason to reject the thesis of "scientific methods only". In chapter three, I stated that both non-scientific psychological investigation and scientific psychological investigation are required for all possible applications of the constraint of psychological realism. In other words, in some applications of the constraint, one must use at least non-scientific methods, and in other employments, one must use at least scientific methods. There are at least two reasons why scientific methods of psychological investigation alone are not sufficient for some applications of the constraint of psychological realism, at least at the present time. The first is that cognitive psychology is still in its infancy. Psychological experiments and theories which are relevant to testing certain accounts for psychological realism may not exist, either because there has not been enough time to develop them or because, at the present stage of inquiry, the required experimental methods have not yet been developed. Secondly, the interests of cognitive psychologists and epistemologists do not always converge.

Hence, although there may be the time and the means to develop some relevant experiments, there may not be sufficient interest on the part of psychologists. Given that, at least at present, scientific psychological investigation is not sufficient alone for some applications of the constraint of psychological realism, it follows that, if we are to use the constraint, we cannot appeal only to scientific methods. Furthermore, given that the non-scientific use of introspection is adequate for at least some psychological purposes, there is no need to restrict employment of the constraint to scientific methods.

If we had a complete scientific psychological theory of the mind we would not need to use introspection. However, since we are currently far short of having such a theory, for practical purposes, if we are to employ the constraint of psychological realism, we must sometimes appeal to introspection. Hence, the former considerations show that, at least for current practical purposes, non-scientific methods of psychological investigation, such as introspection, are required for some applications of the constraint of psychological realism, such as those of Alston and Cohen. In addition, there is good reason to think that, even when the non-scientific use of introspection is dispensable, it is sometimes preferable to scientific methodology. So even if there were no obstacles to the employment of scientific psychological methods, these methods would not always be

preferable. Thus, even in the ideal case, there would be no need to restrict the constraint of psychological realism to scientific methods. Pollock (1986, p. 171) provides a good example of a case in which introspection is preferable to experimental methodology. If we are trying to answer a question concerning a particular skill, such as "What finger do you use to type a 'w'?", it is much easier to determine the answer on the basis of introspection than it would be on the basis of a laboratory experiment. It would be silly to suggest that a firm answer to this question be postponed pending a controlled psychological experiment. To the extent that a question concerning a psychological capacity required for justified belief by an account of justification is similarly more amenable to introspection than to experiment, there would be no need for using scientific methods to answer the question. Perhaps the question concerning the direct voluntary control of belief which Alston considers would fall into this category.

We can conclude from the preceding observations that the non-scientific use of introspection is not only at least sometimes minimally adequate as a method of psychological investigation, but also indispensable, at least at this point in time, for some applications of the constraint of psychological realism. Moreover, introspection is sometimes preferable to scientific methodology for this purpose even if it is not required. Hence, we cannot and need not

restrict the application of the constraint of psychological realism to scientific methods of psychological investigation.

Furthermore, I think it is clear that there are cases in which scientific psychological investigation is required in the employment of this constraint. For instance, Kornblith's scrutiny of the reliability of introspection itself could not be undertaken solely on the basis of introspection, since the reliability of this process would need to be assumed in order for there to be a basis to trust the results of the investigation. (In chapter four, I will discuss the apparent tension created by endorsing both Kornblith's conclusions about introspection and the role of introspection in the application of the constraint of psychological realism.) Also, results in cognitive psychology make it clear that there is a lot going on in our minds to which we have little or no introspective access. If an epistemological account has a psychological component which makes reference to a purportedly non-introspectible capacity of the mind, then introspection would be inadequate as a method of testing the account. So, in some applications of the constraint of psychological realism, scientific methodology is required.

Given this indispensability of scientific psychological investigation for some applications of the constraint of psychological realism, we can dismiss a construal of Pol-

lock's objection which we considered in chapter two. This was the form of his argument obtained by treating 'intuitiveness' in a broad sense which includes appeal to introspection. The argument, read in this way, concluded that the appeal to introspection is always sufficient as a means to determine whether a candidate account of epistemic justification is psychologically realistic. Given the considerations advanced in the previous paragraph, it is doubtful that introspection is always sufficient for this task.

The General Form of the Arguments

The form of each of the evaluators' arguments against each respective epistemological view is roughly the same. The general pattern of reasoning is as follows, although there are variations and differences of emphasis, as we shall see. First, the critic pinpoints an explicit or implicit psychological component of the view he is attacking. Next, he appeals to psychological evidence which indicates that humans do not have the requisite psychological capacity. Finally, he concludes that the epistemological view is completely false or at least in need of revision. In only one case is the denial of skepticism explicit. Since it is a crucial premise in the argument for the constraint of psychological realism, we must assume that it is implicit in the other two arguments which include applications of the constraint of psychological realism. Since Kornblith's

argument is not an application of this constraint, he does not need to reject J-skepticism(H) for his conclusion. Of course, any application of the constraint making use of Kornblith's psychological premises would require this rejection.

Summary

In this section, I have provided an overview of three features of the illustrations which we will be discussing in this chapter and the next. First, I identified the two general categories into which each of the epistemological theses scrutinized falls. These "epistemic responsibility" and "coherence" categories contain many current influential accounts of epistemic justification and similar epistemological theses. Second, I listed the psychological components of the views to which the constraint of psychological realism will be applied. I also pointed out how these components are examples of the three kinds of psychological components incorporated by accounts of epistemic justification: theses about belief-states, belief-processes and the belief-system. Third, I specified a taxonomy of kinds of psychological investigation which can be used when applying the constraint of psychological realism. I argued that both non-scientific and scientific methods are required for the various possible employments of the constraint of psychological realism. Specifically, I argued that we can-

not and need not restrict application of the constraint to scientific psychological investigation. Finally, I sketched the general form which each of the following critical arguments takes. It is now time to look at the illustrations themselves.

Two Appeals to Introspection

Recall that the four applications of the constraint of psychological realism which we will discuss are divided into two groups of two. Alston and Cohen rely primarily on the non-scientific use of introspection for their psychological evidence. Cherniak and Kornblith appeal mostly to experiments and theories in current cognitive psychology for their psychological premises. We will look first at the two employments of non-scientific psychological investigation. In the next chapter, we will continue our detailed survey of these illustrations by examining the two uses of scientific psychological investigation. In each case, I will supplement, comment on and criticize the arguments as necessary.

Alston on Chisholm

In "Concepts of Epistemic Justification" (1985), William Alston considers and criticizes various concepts of epistemic justification as a preparation for his own proposed analysis. He focuses on a passage from Roderick Chisholm's Theory of Knowledge (1977) as a representative of one of the accounts of epistemic justification which he

rejects. He says this of the concept expressed by Chisholm's remarks: "Most epistemologists who have attempted to explicate justification have set out a concept of this sort" (p. 60; As examples, in addition to Chisholm, he refers to Ginet 1975, Wolterstorff 1983 and Naylor 1978. One could add to this list Bonjour (1980, p. 55) who cites Chisholm approvingly in this regard.). The passage to which Alston directs his criticism is the following:

We may assume that every person is subject to a purely intellectual requirement - that of trying his best to bring it about that, for every proposition h that he considers, he accepts h if and only if h is true One way then, of re-expressing the locution "p is more reasonable than q for S at t" is to say this: "S is so situated at t that his intellectual requirement, his responsibility as an intellectual being, is better fulfilled by p than by q" (Chisholm 1977, p. 14).

Alston's Construal of Chisholm's View. In order to facilitate our discussion of Alston's use of the constraint of psychological realism in criticizing the view expressed in the preceding passage and similar positions, we need to make a few harmless terminological changes. These changes will be made primarily to increase the simplicity, clarity, and in some respects the plausibility of Chisholm's remarks. I do not think that they are inconsistent with the spirit of the above quotation. At any rate, Alston's argument would apply to Chisholm's explication as it is originally worded as well. In light of this, we can distill the account into the following:

- (3.1) For any person S, proposition p and time t, S is epistemically justified in believing that p at t if and only if, in believing that p at t, S is not violating any epistemic obligations.

Here we have replaced degrees of reasonableness with reasonableness simpliciter and have substituted the phrase 'epistemic justification' for the word 'reasonableness'. In addition, Chisholm's comments are a paraphrase of comparative reasonableness in terms of an intellectual requirement. In (3.1), 'intellectual requirement' is replaced by 'epistemic obligations' to highlight both the possibility of there being more than one intellectual duty and the specifically epistemological character of these requirements.

Finally, Alston stresses that epistemic justification is most naturally understood, not in terms of obligation, but in terms of the correlative concept of permission (p. 60). He suggests that the most natural model for this type of account of epistemic justification is that of the justification of behavior. To say that someone is justified in doing something is to say that in doing it he is not violating any relevant obligations. One violates an obligation just in case one does something which is not permitted, so if one acts in such a way as to avoid violating an obligation, then one's action is permitted. So by analyzing epistemic justification in terms of 'not violating any epistemic obligations', (3.1) is thereby making justified belief a function

of permission to believe a proposition.

(3.1) needs to be supplemented with a specification of epistemic obligations. Once again, Chisholm's intellectual requirement is that a person must try his best to bring it about that, for every proposition p that he considers, he accepts p if and only if p is true. Now if this is to be practicable, it must amount to something like the following two epistemic obligations, which Alston uses as examples (p. 64):

(3.2) Refrain from believing that p in the absence of adequate evidence;

and

(3.3) Believe that p when confronted with conclusive evidence that p with no sufficient overriding evidence to the contrary.

Alston holds that the most natural way of construing (3.2) and (3.3) is to take obligations to believe or refrain from believing as obligations to believe or refrain from believing here and now. Now refraining from believing can be either passive, as in merely not believing, or active, as in deliberately choosing not to believe. Given Chisholm's emphasis on trying one's best to bring it about that one believes all and only true propositions, it seems most natural to interpret (3.2) as enjoining active refraining. Of course, one could avoid violating an epistemic obligation by passively refraining from believing. However, it seems rea-

sonable to suppose that, at times, one could avoid violating (3.2) only by actively refraining from believing a proposition. For instance, one may already believe the proposition which one is obligated to refrain from believing. Adherence to the obligation in such a case would require one to bring it about that one cease believing the proposition in question. It is this active sense of refraining which Alston apparently has in mind. Given this reading, and the understanding of these obligations as requiring satisfaction here and now, one can be obliged to believe or to refrain from believing only if these mental operations are under one's direct voluntary control.

We are now in a position to see that a natural interpretation of (3.1) contains or presupposes an implicit and unobvious psychological component. According to (3.1), in order for a person to be epistemically justified in holding a given belief, the person must not be violating any epistemic obligations like (3.2) in possessing that belief. Similarly, given an analogous account of being justified in refraining from believing a proposition, in order for a person to be justified in refraining from adopting a given belief, the person must not be in violation of any epistemic obligations like (3.3) in refraining from believing. As we now see, in order to avoid violating obligations of this sort, one must be capable of adopting beliefs and discarding beliefs at will. If human beings do not have these psycho-

logical capacities, then it follows, given the natural reading of (3.2) and (3.3), that J-skepticism(H) is true. So the question before us is this. Are believing and refraining from believing cognitive processes which are under our direct voluntary control? That is, is the natural interpretation of Chisholm's view psychologically realistic?

Alston's Use of Introspective Evidence. So far, Alston's discussion of Chisholm's account has followed the general pattern followed by each of the four illustrations. Alston has explicated Chisholm's account of epistemic justification in such a way as to uncover an implicit psychological component. With respect to the psychological component he says, "I find this assumption of direct voluntary control over belief quite unrealistic" (p. 64; Notice that Alston's use of the word 'unrealistic' here is not the same as our use in the phrase "psychologically unrealistic account of epistemic justification." Alston is applying the word to a psychological thesis, not to an account of epistemic justification (this is true of Cohen and Cherniak as well). I will have more to say about his use of the constraint of psychological realism later.). He then adduces what he regards as "strong reasons for doubting that belief is usually, or perhaps ever, under direct voluntary control" (p. 64). He invites the reader to consider the beliefs we acquire about ourselves and the world around us on the basis of experience. These beliefs are formed on the basis of

"perception, self-consciousness, testimony and simple reasoning based on these data" (p. 64). He concludes that in familiar situations of belief-formation, such as seeing a car coming down the street, "the belief-acquisition mechanism is isolated from the direct influence of the will and under the control of more purely cognitive factors" (p. 65).

This is a strong psychological thesis. Alston believes he has "strong reasons" for it, and yet he does not appeal to any scientific psychological experiments or theories to back it up. Although he does not make the nature of his psychological evidence explicit, it is clear that he is relying on the deliverances of introspection. He is expecting that his readers will be convinced that beliefs are not under direct voluntary control merely on the basis of reflecting on their own experiences of belief-acquisition. This type of argument against the thesis of the direct voluntary control of beliefs is not uncommon in the philosophical literature (see for instance Pojman 1985, pp. 38-47 who has a much more extensive introspection-based argument against the thesis). There does not appear to be a single instance of an argument by a philosopher against this thesis which appeals to results in current cognitive psychology. This is probably because the relevant experimental evidence is simply not available. Furthermore, as I mentioned before, this may be the kind of psychological thesis which, by its very nature, is more amenable to test by means

of introspection than by means of a laboratory experiment. The more one introspects about the coming and going of most, at least, of one's beliefs, the clearer it becomes to one that one is unable directly to control these processes. For instance, with respect to perceptual beliefs, it is relatively easy to see that we have no control over what we believe once we are directly attending to something in our perceptual field. Of course, we can deliberately interrupt the operation of our senses by, for instance, plugging our ears or covering our eyes. However, when such interference is absent, the formation of beliefs is automatic and forced. An apparent exception to this is a case in which we have learned not to trust our senses, such as when we perceive a stick in water to be bent when it is really straight. It might seem that in such cases we can decide to believe that the stick is not bent. However, this is not an exception to the thesis of the automatic formation of belief. The belief that the stick is not bent is surely formed automatically once we have learned that our senses are misleading in such cases.

This argument against the thesis of the direct voluntary control of beliefs which is based on introspection does provide strong grounds for concluding that this psychological capacity is completely lacking in human beings. That is, it constitutes good reason to think that the direct voluntary control of beliefs is psychologically impossible.

Some philosophers have advanced philosophical arguments against this thesis and have concluded that it is logically impossible to decide to believe something (Williams 1973b, pp. 107-108; Swinburne 1981, pp. 25-26). However, such arguments have been largely discredited (see for instance, Winters 1979 and Wolterstorff unpublished, pp. 10-14). Consequently, given the current dearth of relevant experimental evidence, the appeal to introspection is the only adequate basis on which to reject the direct voluntary control thesis. However, before we dismiss it completely, can anything be said in its favor which might counterbalance the introspective evidence? The following type of consideration provides what appears to be the only line of defense for the proponent of this thesis. However, I believe that Alston's response is sufficient to diminish its initial plausibility.

An Attempt to Counterbalance the Introspective Evidence. Alston briefly considers cases in which it is not as clear that believing is automatic. Such cases include the apparent decision to believe a philosophical, religious or scientific thesis in the absence of compelling reasons. It is somewhat plausible to suppose that in such cases belief is directly voluntary. Alston's response to such cases is that "insofar as something is chosen voluntarily it is something other than a belief or abstention from belief" (p. 65). That is, if the will is involved at all in such cases, it involves a cognitive process other than believing, such

as accepting a proposition as a working hypothesis. Alston concludes that, even if such cases were cases of the direct voluntary control of belief, "it is clear that for the vast majority of beliefs nothing like direct voluntary control is involved" (p. 65). He then draws his final conclusion about the natural construal of Chisholm's account of epistemic justification as summarized in (3.1) - (3.3): "(Chisholm's account) could not possibly be a generally applicable concept of epistemic justification" (p. 65).

Alston's Implicit Appeal to the Constraint of Psychological Realism. Alston apparently considers the foregoing considerations as being sufficient to defeat the natural interpretation of Chisholm's position and others like it. However, based upon his explicit remarks alone, the falsity of this construal of Chisholm's account and related views does not follow. So far Alston has (1) located an implicit psychological component or presupposition of one reading of Chisholm's view; (2) argued, on the basis of an appeal to introspection, that humans do not have or rarely have the requisite psychological capacity; and (3) concluded that this interpretation of Chisholm's position expresses a concept which "could not possibly be generally applicable." Although this latter conclusion does indeed follow from Alston's introspective evidence, it does not follow directly from that conclusion that the present construal of Chisholm's account is false. For this we need the premise that

"no account of epistemic justification which is not generally applicable can be the correct account of epistemic justification." This additional premise is the constraint of realism ((3.2) at the beginning of chapter three). Of course, the validity of Alston's argument requires only the constraint of psychological realism, since the reason the natural interpretation of Chisholm's view is not generally applicable is that humans do not have a psychological capacity required by it for justified belief. Furthermore, in order for appeal to this additional premise to be justified, we must read into Alston's discussion the denial of J-skepticism(H). These two additional premises must be what Alston had in mind.

A Psychologically Realistic Alternative Reading of Chisholm. Given this partial reconstruction of Alston's argument, we can see clearly how the constraint of psychological realism plays a key role in his criticism of the natural reading of Chisholm's remarks. Furthermore, it is clear that he appeals to introspection for the psychological evidence he needs. It is quite reasonable to think that this particular application of the constraint of psychological realism has shown that Chisholm's view is false, at least when the relevant epistemic obligations such as (3.2) and (3.3) are given their most natural reading. However, it does not follow that (3.1) is false on every interpretation of it. This use of the constraint of psychological realism

does not so much refute (3.1) and its ilk as much as reveal the need for the revision of such views. What is needed is a more careful statement of these positions which is sensitive to considerations of psychological reality.

Alston himself considers such an alternative reading of (3.1). He follows others in suggesting a revision of views like (3.1) along the following lines. A distinction can be drawn between direct voluntary control and indirect voluntary control of beliefs. Some who adopt this distinction argue that, although believing and disbelieving are not subject to direct voluntary control, they are subject to indirect voluntary control (Price 1954, p. 106; Goldman 1978 p. 222; Kornblith 1982, pp. 252-53; Heil 1983a, p. 753; Alston 1985, p. 65). The idea is that, although we cannot simply will ourselves to believe something, we may be able to put ourselves in a position which we know would be conducive to forming a particular belief or a certain kind of belief. Analogously, we cannot catch a cold at will, but we can refrain from sleeping and eating and we can subject ourselves to cold weather and to sneezing people in order to make it more likely that we will catch a cold (Heil 1983b, p. 359). Given that believing and disbelieving are subject to indirect voluntary control, it may be possible to formulate a psychologically realistic account of epistemic justification in terms of obligations to do what can reasonably be expected of us to influence what we believe or disbe-

lieve. Although such a revision of (3.1) might be promising, I will not pursue it at this juncture, since my purpose here is simply to illustrate Alston's use of the constraint of psychological realism. However, in the section concerning Kornblith's criticism of Cartesian epistemology, we will have occasion to consider a few features of such a view (Alston rejects such a revision of (3.1) on the grounds that it fails to satisfy the constraint of intuitiveness.).

An Assessment of Alston's Use of the Constraint of Psychological Realism. Now that epistemologists have been widely sensitized to the falsity of the thesis of direct voluntary control of belief as the result of efforts by Alston and others, it is easy to think that attacks on accounts like (3.1) which presume that they presuppose this thesis are really attacks on straw men. After all, as we have seen, there is more than one way of construing views like (3.1). For instance, Laurence Bonjour protests, in his recent book The Structure of Empirical Knowledge (1985) that "by appealing to the idea of (epistemic) irresponsibility, I obviously do not mean to suggest that belief is voluntary in any simple way" (p. 230; emphasis mine). Bonjour is most likely responding to attacks on a statement he made (in 1980, p. 55) which, as I previously mentioned, is supportive of Chisholm's remarks, quoted above. A recent example of such a charge leveled against Bonjour can be found in Clarke (1986) who labels Bonjour a proponent of "doxastic volunta-

rism", which, given what Clarke says, is just the thesis of direct voluntary control of belief (Kornblith 1983, p. 34 also accuses Bonjour of endorsing this thesis). In spite of Bonjour's insistence about the "obviousness" of his intended meaning, it is not clear that everyone who has proposed or thought about similar views has been sensitive to the nature and psychological reality of the psychological components these views can be construed as having. It may be that only repeated applications of the constraint of psychological realism to such positions has made it obvious to many that accounts like (3.1) can be interpreted as presupposing the thesis of direct voluntary control of belief and that since this psychological thesis is false, such an interpretation must be false as well. I think the best way of evaluating this particular employment of the constraint of psychological realism to the natural reading of views like (3.1) is to say that, whereas initial such uses of it were arguably not directed at straw men, repeated uses of it probably are. But the fact that we can say this of later uses is in large part due to the success of the initial applications.

In spite of this, continued use of terminology which seems to presuppose the direct voluntary control thesis indicates at least that some philosophers are not yet being careful enough in their formulations. At most, it shows that some are either oblivious to the implications of their remarks or even endorse the questionable psychological the-

sis itself. For instance, consider the following way in which Gilbert Harman, who is otherwise quite sensitive to issues of psychological reality, opens a recent article (1985): "I prefer to speak of belief revision rather than reasoning because the theories I am going to talk about are meant to apply as well to the decision to stop believing something one has previously believed as to the decision to start believing something new" (p. 231, emphasis mine). One can decide to believe or disbelieve something only if one has direct voluntary control over one's beliefs. Unfortunately, there is nothing in Harman's article which indicates whether he intends this to be taken quite loosely or whether he is actually subscribing to the direct control thesis.

This ends our scrutiny of Alston's application of the constraint of psychological realism to an interpretation of Chisholm's comments about epistemic justification in (1977). It is now time to turn to Cohen's use of a version of this constraint in criticizing Lehrer.

Cohen on Lehrer

In "Justification and Truth" (1984), Stewart Cohen examines three kinds of accounts of epistemic justification in terms of what they imply about the connection between epistemic justification and truth. One of the types of views he looks at is the category of coherence theories of justification (pp. 285-288). As a representative of this

approach, he singles out Keith Lehrer's theory of justification in Knowledge (1974). The general outline of Cohen's argument is as follows. First, Cohen provides a simplified but accurate statement of Lehrer's analysis of epistemic justification. Second, he raises an objection to the view which amounts to an application of a version of the constraint of psychological realism. Although Cohen does not appeal to any explicit psychological evidence here, it is clear that an appeal to introspection and informal psychological observation is needed to ground his argument. Next, Cohen considers two ways in which Lehrer and others with similar views might appeal to psychology to avoid the charge of psychological unreality and the resultant skepticism. Upon concluding that neither of these alternatives will work, Cohen considers how Lehrer might resort to a limited skepticism. Although Cohen does not conclude that Lehrer's view is false, he does ascribe to Lehrer the burden of proof for showing that his theory would work.

Cohen's Statement of Lehrer's Account. With this general framework in mind, let us now proceed to examine the details of Cohen's argument. Here is Cohen's statement of Lehrer's analysis of epistemic justification (I have changed the wording slightly to bring it into conformity with our previous terminology):

- (3.4) S is epistemically justified in believing that p if and only if S accepts that, given the set

of propositions which S accepts in the interest of obtaining truth and avoiding error, p has a higher probability of being true than any proposition with which p competes.

Unlike Chisholm's view as expressed in (3.1), (3.4) has an explicit psychological component. This requirement of a psychological capacity for the possession of justified beliefs can be stated more simply. Let T be the set of propositions which S accepts in the interest of obtaining truth and avoiding error. In addition, since one of the competitors of p is not-p, p's having a higher probability of being true than any statement with which p competes entails that it is more probable than not that p is true. Given these two considerations, we can state a consequence of (3.4) which makes its psychological component clearer:

(3.5) S is epistemically justified in believing that p only if S accepts that, given T, it is more probable than not that p is true.

(3.5) makes it clear that, on Lehrer's view, a person can be justified in holding a belief only if the person has another meta-belief about the probability of the first-level belief's being true given other beliefs the person has.

Cohen's Explicit Use of a Version of the Constraint of Psychological Realism. Cohen calls the psychological component that we have isolated an "intellectualist model of justification". He raises an objection to (3.5) which

denies that this intellectualist model is necessary for the epistemic justification of every kind of belief. He uses perceptual beliefs as an example of a kind of belief which can be justified in the absence of such an elaborate system of supporting beliefs. As an alternative approach to Lehrer's account of the justification of perceptual beliefs, Cohen mentions Pollock's (1974, pp. 58-64) contention that "it is sufficient for S to be (at least prima facie) justified in (the) perceptual belief (that there is something red) that S merely be appeared to redly" (p. 286). Cohen also states Lehrer's likely response to this suggestion: "the mere fact that S is appeared to redly cannot justify a belief for S that something is red. S must believe he is appeared to redly and moreover believe that his being appeared to redly makes it likely that there is something red" (p. 286). Would Lehrer be correct in making this response? If not, what kind of consideration could be advanced in order to back up the claim that the intellectualist model is not necessary for justified belief?

Cohen's full objection to (3.5) and hence to (3.4) should be stated in its entirety because of what it reveals about his adherence to a version of the constraint of psychological realism in criticizing Lehrer's position:

Although the claims of the intellectualist have some intuitive appeal - one feels that something like this must be involved in justification - the strongest argument against them seems to be that they run the risk of skepticism. For in many cases of what

we generally take to be e.g., perceptual knowledge, the psychological reality of the supporting beliefs required by the intellectualist model is questionable. There is no obvious sense in which most adult persons, not to mention children and animals, have beliefs about how they are appeared to as well as beliefs about their reliability as perceivers, when they have perceptual beliefs (p. 286).

It is clear from this passage that Cohen is not attacking Lehrer's position by arguing that it is unintuitive. As a matter of fact, he finds it to have "some intuitive appeal". Instead, the primary premise of his counterargument is the rejection of skepticism about justified beliefs. So his argument employs a version of the constraint of psychological realism. Now the reason I have been saying that Cohen uses a version of this constraint is that the kind of skepticism which he believes Lehrer's account to entail, and which he denies, is not J-skepticism(H) as I have defined it. J-skepticism(H) is the view that humans are not capable of possessing epistemically justified beliefs. The skepticism which Cohen denies is the thesis that humans do not actually possess justified beliefs. This latter kind of skepticism is consistent with the falsity of J-skepticism(H). That is, humans can be capable of possessing justified beliefs and yet not actually have any. Given this, Cohen's objection to Lehrer's view does not imply that the latter entails that humans do not have the capacity for justified beliefs. That is, Cohen's argument falls short of showing that humans are incapable of

instantiating the intellectualist model which Lehrer's account requires for the possession of justified beliefs. Consequently, Cohen's argument does not incorporate the constraint of psychological realism as I defined it in chapter two. As defined, this constraint holds that no account of epistemic justification is correct which requires a psychological capacity for justified beliefs which humans do not possess.

Although Cohen's argument does not make use of the constraint of psychological realism as defined in chapter two, it does employ an analogous version of the constraint which is based on equally plausible grounds. The version Cohen uses relies on a denial of the claim that humans do not actually possess epistemically justified beliefs. Now it seems even clearer that humans actually have justified beliefs than that they are capable of having them. As a matter of fact, the thesis that humans do possess justified beliefs was the primary basis on which I argued in chapter two that humans are capable of having them. So, if Cohen's argument against Lehrer's position is sound, then although Lehrer's view may not be psychologically unrealistic as defined in chapter two, it is surely false insofar as it entails skepticism about the actual possession of epistemically justified beliefs. Furthermore, although Cohen does not appear to take this step, I believe that his argument could be supplemented in the following way so as to entail

that Lehrer's account is psychologically realistic in the sense defined in chapter two.

What reason could there be, in the absence of the appropriate psychological evidence, to think that humans are capable of having the kind of meta-beliefs required by Lehrer's account for the possession epistemically justified beliefs, except for the claim that we actually possess these kinds of beliefs? But Lehrer has not adduced any psychological evidence in (1974) for the thesis that humans are capable of instantiating the intellectualist model. Furthermore, if Cohen's psychological claim is correct, then no one actually has the kind of meta-beliefs required by Lehrer's account. Consequently, there appears to be no reason to think that humans are capable of satisfying the requirements for justified beliefs which are incorporated in Lehrer's account of epistemic justification. Thus, given what Cohen argues, there appears to be no reason to think that Lehrer's account in (1974) is psychologically realistic in the sense defined in chapter two. As Cohen himself implies after applying the version of the constraint which he employs, this puts the burden of proof on Lehrer to show that his view is psychologically realistic.

Cohen's Implicit Appeal to Introspection. The psychological claim upon which Cohen bases his conclusion that (3.5) is psychologically unrealistic is the last sentence of the above quotation. Although Cohen does not back

up this psychological thesis with any explicit psychological evidence, it is reasonable to suppose that he, like Alston, is counting on the reader to engage in introspection (and extrapolation to the case of other persons) in order to be convinced that it is true. When he says that "there is no obvious sense in which most adult persons" have the requisite meta-beliefs, it is reasonable to think that what he means by this is that even if there is some basis for thinking we have such beliefs, that basis is not provided by introspection. This way of putting his claim is a way of not putting the entire weight of his objection on introspective evidence. So, unlike Alston, Cohen does not conclude that the view he is attacking is false on the basis of introspection. Instead, he uses this appeal to introspection to shift the burden of proof to Lehrer to come up with psychological evidence which would substantiate the view that (3.5) is psychologically realistic. Furthermore, he provides Lehrer with what he seems to consider are Lehrer's only two alternative ways of doing this.

Strategies for Lehrer to Avoid the Charge of Psychological Unreality and Skepticism. The two strategies which Cohen considers for ways in which Lehrer might try to avoid the charge of psychological unreality and skepticism about the actual possession of epistemically justified beliefs are the following. First, Lehrer could try to show that the relevant supporting meta-beliefs are present unconsciously.

Second, Lehrer could attempt to argue that these beliefs are present dispositionally. In the first category, there are two things which might be meant by 'unconscious belief'. The first is the psychoanalytic notion of "repressed belief" which Cohen quickly discards as being unsuited for Lehrer's purpose. The second way of construing 'unconscious belief' is more promising. This way is to explicate such beliefs as the premises of unconscious inferences. Cohen points out that this is the route taken by Harman in Thought (1973, chapters 7 and 8) in his attempt to show that his own explanatory coherence view is psychologically respectable. The upshot of Cohen's response to this suggestion is that it requires empirical support from cognitive psychology and that the relevant support is currently lacking. So if Lehrer were to take this route, he would need to postpone full endorsement and development of his theory pending further progress in cognitive psychology. At any rate, this avenue is not promising for an immediate solution to his problem. (In Lehrer and Cohen (1983), Lehrer attempts to show that a later but similar version of his account of epistemic justification is psychologically realistic by adopting Fodor's input system/central system dichotomy developed in Fodor (1983). Although this strategy is on the right track, Fodor's thesis about cognitive architecture is, by his own admission, a piece of "speculative psychology." The psychological evidence needed to ground Lehrer's account must be

more than speculative in order for it to constitute a demonstration of the psychological reality of his view.)

The second way that Cohen suggests Lehrer might try to demonstrate that his account of epistemic justification is psychologically realistic is to treat the requisite meta-beliefs as dispositional beliefs. Cohen considers two ways in which beliefs might be thought to be dispositional. The first way involves a belief's having previously been occurrent but presently stored in memory. Cohen rightly dismisses this as a candidate for solving Lehrer's problem. We can supplement Cohen's argument here by observing that a belief that one is now being appeared to in such-and-such a way cannot possibly be dispositional in this sense. The second way in which a belief could be dispositional is for the belief to be such that, "if the subject were to consider the relevant propositions, he would assent to them" (p. 287). This suggestion would have the consequence that, for instance, S is justified in believing that there is something red before S only if, if S were to consider propositions like "I am being appeared to redly" and "My being appeared to redly makes it more probable than not that there is something red before me", S would assent to these latter two propositions.

Cohen's dismissal of this suggestion can be summarized in the following way. First, Cohen makes another implicit appeal to introspection by pointing out that people

rarely if ever have beliefs about the way in which they are appeared to even in this dispositional sense. He backs this up by observing that artists must work at characterizing the way in which they are appeared to. The reasoning behind Cohen's remark here seems to be the following: "One can dispositionally believe that, for instance, one is being appeared to redly only if one's considering (in the relevant sense, that is, by means of introspection) the proposition that one is being appeared to redly would prompt one to assent to that proposition. This latter counterfactual can be true only if one can characterize one's being appeared to redly. We, unlike trained artists, can rarely characterize the ways in which we are appeared to. Consequently, we rarely have dispositional beliefs about the ways in which we are appeared to."

Although the soundness of this argument would block Lehrer from using this route to show that his account is psychologically realistic, Cohen provides a second objection which would also undermine Lehrer's attempt even if we have such dispositional beliefs. The gist of Cohen's objection is that S's merely having a dispositional belief that p at t in the sense defined is not sufficient for p to function as an actual justifying reason for any of S's other beliefs at t. This is because p may be a reason at t for S's belief that q at t without it being the case that S's belief that q at t is based on S's dispositional belief that p at t. That

is, using the terminology of chapter one of this dissertation, S's belief that p may provide counterfactual justification for S's belief that q without providing actual justification. The reason for this is that people can discover reasons for beliefs they have which up until that time were not reasons on which these beliefs were based.

Cohen's example illustrates this point nicely. He asks us to consider a trained meteorologist S who believes that it is going to rain on the basis of what a ouija board tells him (Although Cohen does not mention this, it is clear that he has intentionally chosen an example of a paradigmatically unjustified belief.). In addition, S has neglected to consider the proposition that the present weather conditions indicate the imminence of a rainstorm - a proposition which he would assent to if he were to consider it. So the condition for S's possessing a dispositional belief that the present weather conditions indicate the imminence of a rainstorm are satisfied. However, even though this dispositional belief constitutes an excellent reason for S's believing that it is going to rain, it is clear that S's belief that it is going to rain is unjustified since this belief is based on S's beliefs about the ouija board and not on the previously-described dispositional belief. Hence, even if we always had such dispositional beliefs, their presence could not serve as a sufficient condition for the justification of any of our other beliefs. Consequently,

this route is inadequate for Lehrer's purpose of avoiding psychological unreality and skepticism about the possession of justified beliefs (I will have more to say about counterfactual vs. actual justification in chapter five.).

Skeptical Alternatives for Lehrer. Having exhausted what he considers to be the only routes open to Lehrer for avoiding skepticism about the possession of epistemically justified beliefs as a result of the psychological unreality of (3.5), Cohen considers a way in which Lehrer might reasonably endorse a form of skepticism. This final option for Lehrer is his simply sticking to his account as initially stated and insisting that, when the relevant supporting meta-beliefs for a given belief are absent, the person is not justified in holding that belief. Cohen points out that the degree of skepticism involved in such a stance depends on the answer to the empirical psychological question concerning the extent to which the requisite meta-beliefs are lacking in actual cases. If it should turn out that we never have these kinds of beliefs, then Lehrer would be committed to skepticism about the human possession of epistemically justified beliefs. If instead we sometimes have such beliefs but do not have them every time commonsense says that we have a justified belief, then a more limited form of this kind of skepticism would be the case. Cohen suggests that, as a last resort, Lehrer could view his account of epistemic justification as "an ideal to which as a matter of

empirical fact, we only approximate. . . in some sense, the beliefs are justified insofar as they approximate the ideal model" (p. 288). I doubt that this more limited form of skepticism would follow from Lehrer's view. It seems unlikely to me that, if we have the meta-beliefs required by Lehrer's account for justified beliefs at all, we only sometimes have them. It seems more plausible to say that such beliefs either do not exist or always accompany the perceptual beliefs that we deem justified. If this is the case, then the limited skepticism option is ruled out for Lehrer. This leaves only unlimited skepticism about the possession of justified beliefs as a potential option for Lehrer. However, as I said above, it is quite reasonable to reject this form of skepticism. There are two reasons for this. First, this brand of skepticism seems obviously false. Second, in the absence of any psychological evidence for the thesis that humans have the capacity to have the meta-beliefs required for justified beliefs by Lehrer's account, it seems that skepticism about the possession of justified beliefs would be a good reason to accept J-skepticism(H). But if this is the case, then the former kind of skepticism must be false, since J-skepticism(H) is false. Given these considerations, we can discount any form of skepticism as a possible alternative for Lehrer.

A Comparison of Cohen and Alston. It is now time to sum up our discussion of Cohen's employment of a version of

the constraint of psychological realism in his argument against Lehrer's theory of epistemic justification in Knowledge and to compare his use of this constraint with Alston's. In the first place, Cohen's appeal to a version of the constraint of psychological realism is more explicit than Alston's. Unlike Alston, Cohen acknowledges the key role played by the denial of skepticism in his argument. Secondly, although Cohen and Alston both appeal implicitly to introspective evidence in their application of the constraint, Cohen places less weight on this evidence than does Alston. As we have seen, instead of concluding that Lehrer's view is false on the basis of introspection, Cohen uses introspective evidence as a way of placing the burden of proof on Lehrer to come up with evidence from cognitive psychology which would show that his position is psychologically realistic. This more cautious use of introspection on Cohen's part is appropriate, for the following reason. Since Alston focuses his criticism on an implicit psychological component of a natural, but not the only interpretation of Chisholm's view, he is open to the charge of attacking a straw man. Since it is controversial that the thesis of direct voluntary control of belief is or has been held by any epistemologists, there is not a large presumption in favor of the thesis. Hence, it is appropriate to conclude that it is false on the basis of introspection alone. However, it is quite clear that what Cohen has labeled "the

intellectualist model of justification" is explicit in the views of Lehrer (1974), Harman (1973) and Bonjour (1985). Consequently, there is more presumption in favor of such a view. Since three leading epistemologists endorse this psychological thesis, it would be hasty to conclude that it is false on the basis of introspection alone. In this case, introspection is more appropriately used to shift the burden of proof. However, in the absence of relevant support from cognitive psychology, there is reason to be doubtful about the success of the three views mentioned above and others like them.

Three Coherence Theorists on Skepticism and Psychological Reality. One brief final observation is of interest before we proceed to the next two illustrations in chapter four. I would like to compare Lehrer, Harman and Bonjour with respect to their responses to skepticism and the consequences of these responses. On the one hand, Lehrer (1974) and Harman (1973) are examples of epistemologists who deny J-skepticism(H) but endorse arguably psychologically unrealistic accounts of epistemic justification and knowledge. In chapter three, I appealed to the notion of conditional intuitiveness to explain how this could be the case. Each of these two philosophers believes that the "intellectualist model of justification" is somehow instantiated in human beings. This explains how they can consistently hold their views and deny J-skepticism(H). On the other hand,

Bonjour (1985), like Goldman (1986), leaves the possibility of skepticism concerning epistemic justification open. Consequently, he can and does consistently admit that his epistemological theory may be psychologically unrealistic (pp. 151-152). That is, he would agree that the intellectualist model may not be instantiated in human beings. Of course, a main thesis of this dissertation is that Bonjour would be wrong to endorse a psychologically unrealistic account of epistemic justification or knowledge. Moreover, even if Harman and Lehrer have mistaken epistemological views, they are correct in their denial of J-skepticism(H) and their attempts to provide some reason for thinking that their views are psychologically realistic (see Lehrer and Cohen (1983) and Harman (1973)). Overall, this comparison of these three epistemologists shows that, whereas Lehrer and Harman at least implicitly endorse the constraint of psychological realism, Bonjour explicitly rejects it.

Summary

In this chapter we have seen two illustrations of applications of the constraint of psychological realism. Both of them appeal to psychological evidence collected on the basis of a non-scientific use of introspection. Alston focused on Chisholm's account of epistemic justification which is in the "epistemic responsibility" category of epistemological views. The psychological component of this

position isolated by Alston has to do with a thesis about belief-processes. Cohen expended his critical efforts on an account of epistemic justification in the "coherence" category. This view of Lehrer's has a psychological component which has to do with a thesis about certain kinds of belief-states. We now turn to chapter four, where the remaining two illustrations will be examined. They both make extensive use of scientific cognitive psychology.

CHAPTER FOUR

APPLYING THE CONSTRAINT OF PSYCHOLOGICAL REALISM:

SCIENTIFIC PSYCHOLOGY

In this chapter, we continue our examination of illustrations of the constraint of psychological realism. As I said in chapter three, one of these illustrations employs a different version of the constraint of psychological realism than the one I defended in chapter two. This variation of the constraint is applied by Cherniak. The details of how Cherniak's version differs from ours will be discussed in the section entitled "Cherniak on Quine" below. I also pointed out in chapter three that Kornblith's argument does not make use of any version of the constraint of psychological realism. In spite of this, his discussion is included because his argument has the same general structure as arguments which incorporate the constraint of psychological realism and because it is suggestive of alternative ways in which the constraint might be employed. In addition, Kornblith's thesis that introspection is unreliable for important applications threatens to undermine the arguments of Alston and Cohen. We will need to determine if his arguments for this thesis succeed, and if so, what is implied about the use of introspection in applying the constraint of psychological realism.

After our examination of the remaining two illustrations has ended, we will proceed to an assessment of three general consequences which would follow upon the success of the four applications of the constraint of psychological realism which we will have observed. As I have said, I will argue that these consequences of widespread psychological unreality would underscore the importance of the constraint and the indispensability of psychological investigation in its application. In addition, I will discuss what bearing the success of the evaluators' arguments would have on future research which is aimed at determining the nature of epistemic justification.

Two Appeals to Cognitive Psychology

In this section, we will discuss two applications of the constraint of psychological realism which draw on theories and experiments in current cognitive psychology. First, we will examine Cherniak's criticism of Quine's model of the human belief system. Our final illustration will involve Kornblith's attack on the role of introspection in Cartesian-style epistemology.

Cherniak on Quine

Christopher Cherniak's critique of Quine occurs in the third chapter of his recent book, Minimal Rationality (1986), which is entitled "Rationality and the Structure of Human Memory" (the chapter was previously published with the

same title as Cherniak 1983). Two goals of his chapter are to argue against a Quineian epistemological thesis and to argue for a psychologically realistic thesis concerning what Cherniak calls "minimal rationality". Although I will focus mostly on the first of these two tasks, I will also have a little to say about the second.

The Relevance of Cherniak's Discussion to Our Thesis. One thing should be clarified before we begin our examination of Cherniak's discussion. One might wonder what relevance Cherniak's remarks have to the primary epistemological concern of this dissertation, that is, epistemic justification. One could object that, although rationality is also an epistemological notion, it is not the same as justification, so Cherniak's consideration of the psychological realism of theories of rationality has no relevance to our topic. There is one rationale for including Cherniak's discussion which should be rejected at the outset. However, there are also some good reasons for its inclusion which I will also point out.

The inadequate argument for the relevance of Cherniak's remarks to this chapter is the following: "The above objection fails to take account of an important connection which exists between justification and at least minimal rationality. In order to be capable of having justified beliefs, one must be capable of having beliefs. According to Cherniak, a person cannot have any beliefs unless the

person is at least minimally rational (and the person is minimally rational only if the person maintains some consistency among his supposed beliefs; p. 60). So a person is capable of having justified beliefs only if the person is at least minimally rational. Given this, any psychological requirements for minimal rationality will also be psychological requirements for justified belief. So if a theory of rationality is psychologically unrealistic, that is, if it requires for the possession of rationality a psychological capacity which humans do not possess, then any account of epistemic justification which presupposes it will be psychologically unrealistic as well. Hence, Cherniak's discussion is quite relevant to our primary concern." The main problem with this argument is the questionable premise that "a person cannot have any beliefs unless the person is at least minimally rational." This is a substantive and contentious metaphysical thesis which I am not inclined to endorse. Moreover, it would be beyond the scope of our discussion to consider arguments either for or against this claim. Consequently, I will not assume here that the reason Cherniak's discussion is relevant to the concern of this chapter is that there is a conceptual link between epistemic justification and rationality.

Why then is Cherniak's treatment of Quine included here? There are three reasons which are sufficient to vindicate this inclusion. In the first place, even though I

have argued for a constraint of psychological realism on accounts of epistemic justification, it is quite reasonable to think that there is a strictly analogous constraint which applies to theories of human rationality. Illustrating the use of such an analogous constraint would demonstrate that the central thesis of this dissertation is generalizable to other important epistemological notions. Secondly, if Cherniak is correct in his interpretation of Quine's model of the human belief system, then it is reasonable to suppose that such a model is a component of Quine's program of naturalizing epistemology. Although it is unclear whether Quine has a theory of epistemic justification, it may be that this model is part of what is required for making true beliefs knowledge for Quine. So Cherniak's arguments against Quine could be directed at an epistemological view which is quite close to being an account of epistemic justification. Finally, it is likely that models of memory similar to the one which Cherniak attributes to Quine are featured as components of other recent epistemological positions which are intended as accounts of epistemic justification. These include the views of the three coherence theorists which were mentioned at the end of the last chapter, that is, Harman (1973), Lehrer (1974) and Bonjour (1985). (Goldman (1986, pp. 206-207) concludes that Harman's view of inductive inference in (1973, pp. 158-159) is psychologically unrealistic because it incorporates a model of memory simi-

lar to the one Cherniak attributes to Quine.) So Cherniak's arguments against a psychologically unrealistic model of human memory and against the psychologically unrealistic theories of human rationality which incorporate them could be applied, mutatis mutandis, to accounts of epistemic justification with similar psychological components. Given these three considerations, the relevance of Cherniak's investigation to our main concern is clear. Given this preliminary clarification, we can now go ahead with the illustration.

Cherniak's Interpretation of Quine. Cherniak begins his critique of Quine's model of human memory by providing what he considers to be a natural interpretation of a passage from "Two Dogmas of Empiricism" (Quine 1961). Here is the passage:

Total science is like a field of force whose boundary conditions are experience. A conflict with experience at the periphery occasions readjustments in the interior of the field Reevaluation of some statements entails reevaluation of others, because of their logical interconnections . . . (p. 42).

Here is Cherniak's interpretation of this quotation:

A natural "descriptive" interpretation here is that Quine is claiming that reevaluation of one belief entails reevaluation of others to maintain consistency of the total system, and that he is predicting that the proprietor of a belief system will in fact make the appropriate reevaluations. Given Quine's naturalism, one would expect that his account is intended to be descriptively correct (p. 50).

Now I am not as concerned here about whether Cherniak has

gotten Quine's view right as I am about how he criticizes Quine's remarks as he interprets them. Of course, if Cherniak is wrong about what Quine means here, then he may not have shown that Quine holds a psychologically unrealistic view of the human belief system and a psychologically unrealistic epistemological position. However, it is reasonable to think that something like Cherniak's interpretation of Quine is presupposed in many philosophical discussions, even if not in Quine's writings themselves. Although I am not committed to Cherniak's reading of Quine's remarks, for simplicity I will speak of Cherniak's construal of Quine as "Quine's model of human memory" (or a similar locution) in what follows.

Cherniak's Appeal to Scientific and Commonsense

Psychological Theories. One way in which Cherniak refers to this model is as "Quine's model of consistency maintenance" (p. 56). He argues that this model is psychologically unrealistic because it "does not take into account the basic organization of human memory" (p. 51). Given that this model is a component of Quine's theory of knowledge, it follows that Quine's epistemological view is psychologically unrealistic as well. What is the basis on which Cherniak dismisses Quine's model of consistency maintenance as being psychologically unrealistic? Interestingly, Cherniak refers to both scientific and non-scientific psychological theories to support this conclusion. First, he sketches three differ-

ent types of scientific theories of human memory. It is clear that, if any of these types of theories is even approximately correct, then neither Quine's model nor more realistic versions of it are instantiated in human beings. Second, Cherniak argues that there is a commonsense psychological theory of human memory which shares at least two components with each of the three scientific psychological theories he has outlined. Given only these two theses, Quine's thesis about human memory cannot be right.

Let us take a look at some of the details of these theories of memory. The three types of scientific theories of human memory which Cherniak discusses are (1) the standard model of the verbal learning and memory tradition (hereafter, the "standard model"); (2) "semantic memory" theories; and (3) "constructive memory" theories. Cherniak describes these types of theories, as well as our commonsense psychological theory of memory, as sharing one general model of human memory (p. 52). So these theories are just variations on that one theme. Cherniak sketches most of the basic structure of this general model in his discussion of the standard model. In his brief treatment of the latter two types of theories, he pauses only to point out a few ways in which they vary from the first type of theory.

The most fundamental feature of the standard model is that it has a "duplex" structure. That is, it divides the human memory system into long-term memory (LTM) and

short-term or active memory (STM). Although there may be many special purpose LTMs and STMs, there are only these two kinds of memory on this model. These two kinds of memory can be distinguished in terms of how they differ with respect to function, capacity, duration, and degree of processing. In addition to mentioning these differences, I will discuss an important processing relation which holds between them and talk about an additional important feature of LTM. The overview of this model is incomplete, but sufficient for present purposes.

The main function of STM is to provide a locus for current thought. The contents of STM are the things a person is thinking about at a given time, whether consciously or unconsciously (as when one attends to driving a car while talking). The capacity of STM is quite limited. Roughly seven meaningful units of information can be contained in STM at a given time (Miller 1956). So whatever is in STM at a given time comprises a very small subset of a person's beliefs and other remembered cognitive states. In addition, these few items will leave STM after about half a minute unless they are rehearsed by the subject. So the duration of storage in STM is limited as well. Finally, with respect to processing, the standard model holds that all theoretical and practical reasoning takes place in STM. Nothing can affect thought or behavior unless it is processed in STM.

The primary role of LTM is to be a repository for

the massive number of cognitive states which are stored in a person's mind but not currently in STM. There is virtually no practical limit to the number of cognitive units which can be stored in LTM. Moreover, the length of time during which these mental items can be stored in LTM is indefinite. Again, it is clear that most of a person's cognitive system is contained in LTM at any given time. On the standard model, processing in LTM is non-existent. No changes can be made to any cognitive state unless the state first enters STM.

This brings us to the most important relations between LTM and STM. In addition to the initial encoding of an item into LTM which proceeds via STM, encoded items can be "copied" from LTM to STM by means of retrieval. At this juncture it will be fruitful to supplement Cherniak's summary of the standard model with some recent observations about retrieval. There is psychological evidence which supports the hypothesis that memory failures are due in large part, not to problems at the encoding stage, or to memory decay, but to failures to retrieve encoded information (Anderson 1984, pp. 5-6). Two main factors have been found to affect the probability of recall. The first is that it is easier to recall something from memory when the context of recall is similar to the context in which the item was encoded (Tulving 1983). The second is that retrieval can be hampered by certain interference effects. Information about

a concept can resist recall when additional information about the concept is acquired. Tests indicate that the prior information is still in LTM but that access to it has been destroyed by the added information (Wickelgren 1976). These considerations seriously threaten the thesis that any mental state stored in LTM can be retrieved into STM at any time, a thesis implied by the Quineian model. Access to items stored in LTM is contingent on at least context similarity and the presence or absence of interference effects.

The final feature of the standard model which Cherniak mentions is that the contents of LTM are organized. An item stored in LTM is located for retrieval into STM, not by means of an exhaustive search of LTM, but by means of a more limited search which exploits the organization of LTM. Cherniak describes LTM as a "graph-theoretic" entity. The intersections of the graph lines are nodes which contain bundles of stored information in much the same way that a file contains subfiles. The lines connecting these points are "arcs" which serve as the pathways for memory searches. On the traditional view, these arcs are associative links. Cherniak's remarks about the organization of LTM suggest that he would agree with Anderson (1984, p. 6) that "it makes . . . sense to view the human's memory difficulties as arising from . . . the problem of how to achieve rapid and flexible access to all of this knowledge (stored in LTM)." The organization of LTM appears to be at least a partial

solution to this problem.

The two additional kinds of accounts of memory which Cherniak mentions vary in specific ways from the standard model but the general features are much the same. "Semantic accounts" have arisen from the computer modeling of human memory. Two variations on the standard model which they incorporate are (i) a generalization of the interconnecting arcs in LTM to represent different kinds of relations between nodes (such as 'is a subset of') and (ii) a view of the structure of LTM as being dynamic. This latter property is exhibited in the co-existence of different organizational systems and the possibility of restructuring memory for different purposes. According to "constructive accounts" what is encoded in LTM is not a copy of the experience it represents. Accordingly, recall involves a reconstruction of "what must have occurred". The main way in which constructive accounts differ from the standard model is that there is a certain amount of processing of the contents of LTM. New information can be assimilated to previously stored items. In spite of this, processing in LTM is still quite a bit more limited than processing in STM.

This ends the review of three important types of scientific models of human memory. Cherniak adds that there is a commonsense model of human memory which, although being less sophisticated than the scientific models, agrees with them on two salient points. The details of his discussion

of this commonsense theory need not concern us, since we are focusing here on his appeal to scientific methodology in applying the constraint of psychological realism. It will be sufficient simply to mention the two theses shared by both scientific and commonsense theories. First, Cherniak argues that the commonsense theory also recognizes the duplex structure of memory. That is, it acknowledges the existence of LTM and of STM. Secondly, the non-scientific theory assumes that LTM is organized. Cherniak goes on to show how these two assumptions figure in commonsense psychological explanations of lapses from ideal rationality.

The Psychological Unreality of the Quineian Model.

We now have sufficient psychological evidence to conclude that the Quineian model of memory is psychologically unrealistic. Cherniak rephrases Quine's model as an information-processing model to make it somewhat more realistic. In terms of this type of model ". . . Quine might be described as viewing the entire belief system as contemporaneously fully activated or processed in parallel; the contents of the short-term working memory would be the complete long-term memory" (p. 70). It is now easy to see that even this amendment is psychologically unrealistic. Even if constructive accounts of memory are correct, so that there is some processing of mental states in LTM, it is clear that LTM cannot be "contemporaneously fully activated or processed in parallel." Furthermore, this revised model would

fail to recognize the basic duplex structure of memory. Cherniak considers a somewhat more realistic Quineian model: ". . . although short-term memory capacity is limited, all and only the appropriate beliefs at a given moment are always recalled" (p. 70). Obviously, given the above psychological theories, this model is psychologically unrealistic as well. This model fails to take account of the way in which LTM is organized. Although it is reasonable to think that LTM is organized in such a way as to facilitate the retrieval of relevant information, the persistent failure of humans always to recall every stored item relevant to an inference or decision is ample evidence against the above thesis. So even a more psychologically realistic version of the Quineian model does not succeed.

Does Cherniak adduce more psychological evidence than necessary to draw the preceding conclusion? Perhaps, but there are two reasons for being as complete as one can in investigating a psychological thesis. The first is that one may need to evaluate successive attempts to make such a thesis more psychologically realistic, as Cherniak has done above. The more psychologically realistic a thesis becomes, the more the need to consult theoretical details to test the thesis for complete psychological reality. The second reason for thoroughness in consulting psychological evidence is that, the more one knows about the human mind, the more one is guided in constructing and testing a psychologically

realistic epistemological account. Cherniak makes use of the psychological theories he has sketched in this latter way as well, as we shall see.

Cherniak's Employment of a Version of the Constraint of Psychological Realism. In addition to concluding that Quine's psychological thesis is psychologically unrealistic, Cherniak argues that no conception of rationality which requires such an unrealistic belief system for the possession of rationality is adequate for employment in the explanation of "important ranges of actual human behavior" (p. 49). Cherniak refers to the kind of rationality which requires the memory capacities of Quine's model as "ideal rationality". The general outline of Cherniak's argument against the adequacy of ideal rationality for psychological explanation is as follows: "The notion of ideal rationality can be employed in the explanation of human behavior only if humans are ideally rational. Humans are ideally rational only if they satisfy the requirements of ideal rationality. They cannot satisfy these requirements unless they are able to make all useful inferences which are relevant to the actions they perform. In order to make these inferences, they must have unlimited resources of memory and time (this is a feature of the Quineian model of memory). But, according to our best psychological models of human memory (which were sketched above), we have less-than-unlimited resources of memory and time. Consequently, the notion of ideal

rationality cannot be employed in the explanation of human behavior."

Now one possible response to this argument is as follows. Suppose that a person is rational only if the person is ideally rational. Given this, and the above argument, it would follow that no human being is rational and that, since a person's behavior is explicable only if the person is rational, no human's behavior can be explained. This is analogous to our concluding that J-skepticism(H) is true on the basis of deciding that some psychologically unrealistic account of epistemic justification is correct. However, Cherniak holds, and I agree, that human behavior is explicable and that, consequently, humans are rational (pp. 54-55). Since humans are not ideally rational, they must be less-than-ideally rational. Another way of putting this is to say that some psychologically realistic standard of rationality applies to them. As Cherniak puts it, "an adequate understanding of the rationality of an agent's actions is not possible without (psychologically) realistic (models of human memory)" (p. 49). So Cherniak would deny the premise of the preceding argument according to which a person is rational only if the person is ideally rational. Human beings can be rational without satisfying the requirements of ideal rationality.

We are now in a position to see how Cherniak makes use of a different version of the constraint of psychologi-

cal realism than the one employed by the previous two evaluators. Although Cherniak's discussion closely parallels those of Alston and Cohen with respect to the general structure of their arguments which incorporate the constraint of psychological realism, there is one important disanalogy between his argument and those of the others (other than the fact that he focuses on the rationality of behavior instead of the justification of beliefs). In applying the constraint of psychological realism to the accounts of epistemic justification which they are criticizing, Alston and Cohen each assume that if the view they are attacking is psychologically unrealistic, then it is false. Cherniak, however, does not conclude that theories of ideal rationality are false. Instead, he concludes that they do not apply to human beings. What is the reason for this difference between these evaluators? The explanation for this difference is that, whereas Alston and Cohen are each criticizing what purports to be a general account of the nature of epistemic justification, Cherniak is instead focusing on a standard of rational behavior (refer to chapter one for a discussion of the difference between general accounts of epistemic justification and standards of epistemic justification). A general account of the nature of an epistemological property such as justification or rationality can be true or false. However, various standards according to which a belief is justified or a person is rational cannot

be true or false. Of course it can be true or false whether or not a given standard applies to a given type of person. This explains why Cherniak, who is talking about standards of rationality, concludes not that theories of ideal rationality are false, but that they do not apply to human beings. Given this, however, there is a short route to the conclusion that a general account of rationality cannot incorporate the requirements of ideal rationality. If it did, then, given a constraint of psychological realism on accounts of rationality, such an account of rationality would be false. This is because the account would be psychologically unrealistic. We can conclude that, although Cherniak applies a distinct form of the constraint of psychological realism according to which no psychologically unrealistic standard of rationality applies to humans, it is a trivial matter to convert his conclusion about Quineian theories of rationality to one formally equivalent to the conclusions of the other critics.

Cherniak's Argument for a Psychologically Realistic Thesis about Minimal Rationality. As I previously mentioned, Cherniak not only criticizes Quine's model of memory and theories of rationality based on it in (1986), he also argues for a positive thesis about minimal rationality which he claims incorporates a psychologically realistic model of memory. Very briefly, his argument for this positive thesis is as follows: "Minimal rationality requires efficient

recall, since one cannot act minimally rationally unless one has access to at least some of one's beliefs which are relevant to the action and such access requires efficient search and retrieval of items in LTM. Efficient recall requires incomplete search, since a complete or exhaustive search of LTM would be quite costly in terms of time and mental resources. Incomplete search requires the compartmentalization of LTM, since incomplete searches could not succeed if the contents of LTM were unorganized and random and the avoidance of such arbitrary storage requires a degree of compartmentalization. Hence, minimal rationality requires some compartmentalization of the contents of LTM."

Cherniak on Psychological Realism and Psychological Investigation. Cherniak claims that the thesis argued for in the previous paragraph is "not just an accident of human psychology Given the short-term/long-term memory distinction and some other basic constraints, the long-term memory could not be otherwise, or the total belief system could not maintain minimal rationality" (p. 61). In addition, with respect to the separate thesis that minimal rationality requires the attempt to maintain consistency in one's belief set, Cherniak claims that "this is one of the minimal rationality conditions on any belief system; we can be sure it applies before we try any experiment" (p. 60). These two assessments of theses about minimal rationality which Cherniak endorses suggest that he would deny the

indispensability thesis which I argued for in chapter two. That is, Cherniak appears not to agree that psychological investigation is required to test an epistemological thesis for psychological reality. Of course in the above quotation he states that we can determine that a minimal rationality condition applies without engaging in an experiment. Presumably, what he means by 'experiment' is 'scientific psychological experiment.' This seems to leave open the possibility that he would hold that theses about minimal rationality must be tested for psychological reality by introspection or some other non-scientific method. However, at no point does he appeal to such non-scientific methods in his argument for the "compartmentalization thesis". Neither does he suggest that one must introspect to determine if one's LTM is compartmentalized. His argument for this thesis about minimal rationality proceeds entirely on the basis of a priori thought experiments about how the mind must be constructed given both the STM/LTM distinction and the thesis that humans are minimally rational. Consequently, Cherniak accepts, as being both psychologically realistic and true, a thesis about minimal rationality which has a psychological component, on the basis of sheerly a priori considerations.

Given this, we can spell out more specifically how Cherniak's approach differs from ours. Consider the following argument: "Humans are minimally rational. Minimal

rationality requires the compartmentalization of LTM. Consequently, LTM is compartmentalized." I am inclined to agree with Cherniak that this argument is sound. However, Cherniak and I would disagree about the basis upon which to determine that it is sound. I claim that the second premise cannot be known without first knowing that the conclusion is true. Furthermore, my contention is that the conclusion cannot be known apart from psychological investigation (such as the psychological investigation which lead to the development of the theories of memory summarized above). So the second premise cannot be known to be true apart from psychological investigation. This is an instance of the indispensability thesis which I defended in chapter two. My claim about the above argument is tantamount to saying that we cannot demonstrate a psychological thesis by some indirect means which bypasses direct psychological investigation concerning that thesis. Cherniak appears to agree with Pollock that one can know that the second premise is true without knowing that the conclusion is true. That is, he would say that one can accept the epistemological thesis as being psychologically realistic apart from psychological investigation. This is tantamount to claiming that one can demonstrate a psychological thesis without engaging in direct psychological investigation concerning that thesis. In sum, although Cherniak does employ a version of the constraint of psychological realism in his argument against Quine, his

apparent denial of the indispensability thesis reveals that he considers his use of psychological theories to be unnecessary. Of course he does need the STM/LTM distinction to serve as a starting point for his argument for the "compartmentalization thesis."

Summary. In this section, we have seen how Cherniak applies the constraint of psychological realism to his interpretation of some remarks of Quine concerning the human belief system. In this employment of the constraint, Cherniak appeals to both scientific and non-scientific theories to back up his claim that Quine's model of human memory and theories of rationality which incorporate it are psychologically unrealistic. Cherniak also argues for a thesis about minimal rationality which he claims is psychologically realistic. It appears that he would draw this conclusion about the thesis even in the absence of relevant psychological investigation. This suggests that, although he views such investigation as sufficient for testing epistemological theories for psychological reality, he does not think that it is necessary. So, although Cherniak endorses a constraint of psychological realism, he diverges from our thesis that the constraint requires appeal to psychological investigation. In the next section, we will see how Kornblith employs an argument which is similar in general structure to arguments incorporating the constraint of psychological realism.

Kornblith on Cartesian Epistemology

In "Introspection and Misdirection" (Unpublished), Hilary Kornblith critically examines an epistemological thesis which he attributes to Descartes and to current epistemology which is done in the Cartesian tradition. As was the case with the epistemological theses criticized by the previous three evaluators, the truth of this Cartesian thesis depends on the truth of a certain psychological thesis. However, unlike the views targeted by those three critics, the object of Kornblith's criticism is not a thesis about a fundamental epistemological notion such as knowledge, epistemic justification or rationality. Consequently, Kornblith's argument is not an instance of the constraint of psychological realism as I have defined it (that is, as applying to accounts of epistemic justification), nor is it a generalized version of that constraint so-defined. In spite of this, there are good reasons for including Kornblith's discussion at this juncture. After the following discussion of Kornblith's remarks, I will comment on the relevance of his arguments to the present project.

The Epistemological Thesis Which Kornblith Denies.

What epistemological thesis does Kornblith attribute to Descartes and to contemporary Cartesian epistemology? As was the case with Cherniak's treatment of Quine, I am not as concerned about Kornblith's reading of Descartes as I am

about how Kornblith argues against a thesis that he believes, rightly or wrongly, Descartes to have espoused. Kornblith attributes to Descartes the view that ". . . the primary task of the epistemologist (is) the giving of advice to those concerned with the search after truth" (p. 1). What general piece of epistemic advice does Kornblith think that Descartes was concerned to give? A second Cartesian thesis provides some direction here. Kornblith ascribes to Descartes the thesis that ". . . the mind is better and more easily known than the external world" (p. 1). Given this, it would be advisable to start one's quest for truth by means of coming to know what is going on in one's mind. This knowledge of one's mental life will then serve as a foundation for the rest of what one knows. But what is the best method for gaining knowledge about one's mind? Here the Cartesian, according to Kornblith, advises one to engage in introspection. We are now in a position to state the general epistemological thesis which Kornblith considers to be at the center of the epistemological theories of both Descartes and contemporary Cartesians: "One's pursuit of truth ought to begin with introspection." This thesis takes a different form for contemporary Cartesians than it does for Descartes. Kornblith phrases Descartes' version as follows: "If one wants to know the external world one ought to begin by introspecting" (p. 2). The contemporary Cartesian version is as follows: "If one wants to improve epistemi-

cally, one ought to begin by introspecting." It is this latter thesis which to which Kornblith directs his criticism. What epistemic component does Kornblith consider this thesis to have? In order to determine this, we need to look more closely at epistemic improvement and a related notion, epistemic responsibility.

The Psychological Component of the Cartesian Thesis.

Kornblith assumes that the most general advice which Cartesians give to truth-seekers is to improve epistemically. Furthermore, in order to improve epistemically one needs to be epistemically responsible. What are the requirements for epistemic responsibility? Kornblith begins his sketch of some necessary conditions of epistemic responsibility by pointing out that the thesis of direct voluntary control of belief (which was discussed by Alston, above) which he attributes to Descartes, is not widely held today. He ascribes to the contemporary Cartesian the view that beliefs are subject to indirect voluntary control (this is the view which Alston proposed after dismissing the other thesis). Given that our beliefs are produced in us ". . . without our (direct) intervention or reflection . . . , the truth-seeking agent ought to be concerned about the nature of these belief-forming processes" (p. 3). The epistemically responsible agent will continually engage in the processes of self-evaluation and self-correction. In particular, such a person ought to determine what his belief-forming pro-

cesses are and whether they are reliable. If they are reliable, he will leave them alone. If they are not reliable he will seek to modify or eliminate them. In addition, he will try to discover additional reliable belief-forming processes which he doesn't have and will try to acquire them. By being epistemically responsible in this way, it is hoped that one will improve epistemically. That is, one will hopefully become more generally reliable as a truth-seeker. In each of the above procedures, the Cartesian considers introspection to play a crucial role. However, it is clear that introspection will contribute to epistemic improvement in this way only if it is reliable. For one thing, unless a person is capable of reliably introspecting which particular processes are responsible for producing his beliefs, he will not be able to determine whether these processes are reliable. Furthermore, unless he can determine whether these processes are reliable, he will not be able to monitor his epistemic improvement. If he cannot monitor his improvement, it is unlikely that he will improve. Consequently, the contemporary Cartesian epistemological thesis has a psychological requirement which can be expressed as follows: "Epistemic improvement ought to rely heavily on introspection only if we can reliably introspect what belief-forming processes we possess and whether they are reliable."

Is Kornblith Attacking His Own View of Justifica-

tion?. Before we go on to consider Kornblith's argument against the Cartesian epistemological thesis stated above, it is important to engage in some parenthetical clarification to forestall a possible misunderstanding which might arise with respect to the relationship between Kornblith's own view of epistemic justification and his present argument. What needs to be clarified are the connections Kornblith makes between the notions of epistemic justification, epistemic responsibility and epistemic improvement. Although Kornblith does not make an explicit connection between epistemic responsibility and epistemic justification in "Introspection and Misdirection", it is clear from another recent article by him that he endorses a conceptual link between these two notions. In "Justified Belief and Epistemically Responsible Action" (1983), Kornblith states his view of epistemic justification schematically as follows:

An epistemically responsible agent desires to have true beliefs, and thus desires to have his beliefs produced by processes which lead to true beliefs; his actions are guided by these desires. Sometimes when we ask whether an agent's belief is justified what we mean to ask is whether the belief is the product of epistemically responsible action, i.e., the product of action an epistemically responsible agent might have taken (p. 34).

According to Kornblith in (1983), then, a person's belief is justified only if that belief is the product of epistemically responsible action on the part of that person. It might be thought that, given what was said above about epistemic responsibility, this latter notion requires a

reliable method of introspection. If this were the case, then Kornblith's arguments against the reliability of introspection would entail that his account of epistemic justification is psychologically unrealistic. However, this line of reasoning rests on a confusion between the notions of epistemic responsibility and epistemic improvement. As defined, an epistemically responsible agent is one who "desires to have his beliefs produced by processes which lead to true beliefs" and acts upon this desire. It is clear that, although acting on this desire may lead one to engage in introspection concerning one's belief processes, one's employments of introspection need not be reliable in order for one to be epistemically responsible. One need only either believe that one's applications are reliable or at least fail to believe that they are unreliable. So epistemic responsibility is largely a subjective notion. On the other hand, epistemic improvement is clearly an objective notion. In order for epistemic self-evaluation and self-correction to lead to true epistemic improvement, the method used for evaluation and correction must be a reliable method. The best way to improve epistemically is to be epistemically responsible. But epistemically responsible use of introspection contributes to epistemic improvement only if the use of introspection is reliable. In sum, epistemic responsibility does not require a reliable method of introspection. Epistemic improvement does. But it is epis-

temic responsibility which is required for justified belief, not epistemic improvement, so Kornblith's arguments against the reliability of introspection do not entail that his account of epistemic justification is psychologically unrealistic. We now return to an examination of Kornblith's argument against the Cartesian thesis.

Kornblith's Main Argument Against the Cartesian Epistemological Thesis. Kornblith's argument has two premises. The first one is a statement of the psychological requirement for the truth of the Cartesian thesis: (1) Epistemic improvement ought to rely heavily on introspection only if we can reliably introspect what belief-forming processes we possess and whether they are reliable. We have already found ample reason for thinking that this premise is true in the discussion concerning epistemic improvement. Kornblith's second premise is a psychological thesis: (2) Introspection is not only unreliable for the purposes stated in (1), it is also misleading. The next section will contain an extended review of his arguments for this claim. The conclusion is, of course, that: (3) The Cartesian thesis is false; it is not the case that epistemic improvement ought to rely heavily on introspection. If epistemic improvement is possible, some alternative, reliable method must replace introspection. We will see at the end of this section what Kornblith's suggestion is for such a replacement method. In addition, at that time we will also provide

an overall assessment of Kornblith's case for (2).

The second premise of Kornblith's main argument is that introspection is both unreliable and misleading. What psychological evidence does Kornblith advance to support this conclusion about introspection? Unlike Cherniak, Kornblith does not appeal to psychological theories to support his psychological thesis. Instead, he rests his case on five recent experiments in cognitive psychology and on one more informal psychological observation. Kornblith uses four of these psychological considerations to argue that introspection is powerless to locate the processes which are responsible for producing our beliefs. He employs the other two experiments to argue that introspection is not only unreliable but leads us to adopt an inaccurate and highly exaggerated opinion of the reliability of our processes. Let us examine each of these psychological considerations and Kornblith's interpretation of them.

Kornblith's Case for Premise Two: Experiment #1.

According to Kornblith, the first two experiments which he examines demonstrate the failure of introspection to detect factors which have influenced a person's judgments. The first experiment is reported by Nisbett and Ross (1980). Here is the passage to which Kornblith appeals to draw the preceding conclusion about introspection:

. . . passers-by in a shopping mall were invited to examine an array of consumer goods (four nightgowns in one study, four identical nylon pantyhose in

another) and to rate their quality. There was a pronounced position effect on evaluations, such that the right-most garments were heavily preferred to the left-most garments. When questioned about the effect of the garments' positions on their choices, virtually all subjects denied such an influence (usually with a tone of annoyance or of concern for the experimenter's sanity (p. 207)).

Kornblith's first conclusion about this experiment is that the subjects ". . . are unaware that their judgments are affected by the relative position of the goods being evaluated" (p. 5). This claim does not yet threaten the reliability of introspection. If it demonstrates anything about introspection at all, it is that the subjects' judgments are not automatically accompanied by introspection of all of the factors influencing their judgments.

Given this first conclusion of Kornblith about the "position effect" experiment, one thing that he says about these cases is puzzling (he says the same thing about the next experiment we will look at as well). He remarks that "the position effect presents a strong prima facie case of faulty inference" (p. 5). He points this out because he hopes to show that introspection is powerless to uncover faulty inferences. However, I fail to see wherein the faulty inference lies. There appear to be only two ways in which it might be claimed that these cases involve inference. The first suggestion does identify a faulty inference, but it is obvious that the subjects do not make this kind of inference. The second suggestion locates a type of

inference which the subjects are probably making. However, it is implausible to say that it is a faulty type of inference.

Here is why the first suggestion fails. Barring unconscious inference, it is clear that the subjects are not inferring that a certain item is superior on the basis of believing that the item occupies a certain relative position. How could it be that they are making such an inference if they are unaware of the role that relative position plays in their judgments? It appears that Kornblith may be conflating the distinction between part of the cause of a belief and reasons for that belief. Even if all the reasons for a belief are at least part of the cause of that belief, it is not the case that all the causal factors involved in the production or maintenance of a belief are reasons for that belief. The relative position of the items in the experiment may have been a causal factor in the production of the subjects' beliefs about the items and yet not have been part of the reasons upon which their judgments about the relative quality of the items were based.

The second suggestion fails for the following reason. If the subjects are making an inference at all, they are inferring that a certain item is superior to the others on the basis of perceived (but illusory) differences in the items. But this is not a faulty inference. It is a legitimate inference based on a false belief. This false belief

may have been caused at least partially by the position of the items, but as I argued in the previous paragraph, this causal relation is not plausibly viewed as an inference. We can conclude that the position effect does not constitute an example of faulty inference. So Kornblith cannot legitimately conclude on the basis of this experimental evidence that reasoners are unable to introspect faulty inferences which they have made.

Kornblith's additional observation about this experiment does, however, cast doubt on the reliability of introspection in detecting all of the factors which influence the formation of beliefs, at least in situations similar to the experimental situation in which the position effect was observed. Even if the judgments of these subjects are not automatically accompanied by introspective access to all of the factors contributing to the formation of their beliefs, one might think that it could still be the case that subsequent careful and deliberate use of introspection would be sufficient to reveal these factors. However, it is clear from the passage quoted above that when the experimenters asked the subjects if the position of the garments had any affect on their judgments, virtually all the subjects denied such influence. Although it is not clear how much introspection these subjects engaged in to draw this conclusion, it is reasonable to suppose that they did introspect to some degree. Kornblith's assessment of

the subjects' denials is that they did introspect and that introspection failed to reveal the affect of position on their choices. He goes so far as to say that even the most epistemically responsible reasoner would be subject to this illusion when introspecting. So Kornblith concludes that even the most careful use of introspection would be to no avail here.

I think Kornblith's case can be strengthened a bit. It seems plausible to suppose that the only way one could become convinced that one has fallen victim to the position effect is by means of inference to this conclusion on the basis of psychological data such as the data collected in the shopping mall experiment. After all, if we could discover such influences on our judgments by means of introspection, it seems likely that phenomena such as the position effect would be less prevalent than they appear to be. At any rate, it is reasonable to conclude from the quoted passage that there is some reason to doubt that we are capable of achieving introspective access to all of the causal factors contributing to the formation of our beliefs, at least in cases of this kind. So far, this thesis falls far short of supporting one of Kornblith's main psychological theses that ". . . our processes of belief-acquisition, and indeed mental processes in general, are largely unavailable to introspection" (p. 14). Does the other psychological evidence which Kornblith adduces make up for this lack of sup-

port? To find out, let us move on to the other experiments.

Kornblith's Case for Premise Two: Experiment #2.

The second psychological experiment which Kornblith employs in his defense of the second premise of the argument against the Cartesian epistemological thesis under consideration is one conducted by Tversky and Kahneman (Kahneman, Slovic and Tversky 1982, p. 14). Kornblith considers this experiment to be another illustration of the thesis that introspection fails to detect factors which have influenced a person's judgments. Tversky and Kahneman call the type of judgment-influencing factor which appears to be operative in the experiment the "anchoring effect". Here is Kornblith's summary of the experiment and its results:

Subjects were asked the percentage of African members of the United Nations. A roulette wheel was spun, and subjects were asked whether the number which turned up was too high or too low. In those cases where the roulette wheel provided an anchor of 10, the mean estimate was 25: where the anchor was 65, the mean estimate was 45. It is quite clear that subjects were influenced in their judgment by the result of the spin of the roulette wheel. None of the subjects, however, showed any awareness of this influence (p. 6).

(It should be noted in addition that, after judging whether the number indicated by the roulette wheel was too high or too low, subjects were asked to estimate the percentage by moving upward or downward from the given number.) Kornblith's conclusion about the reliability of introspection on the basis of this experiment is that ". . . there (is no) reason to think that asking the subjects to introspect more

carefully here would provide much assistance" (p. 7). Now although Kornblith may very well be correct in this assessment, it is by no means clear that it is warranted by his summary of the experiment. Unlike the previous case, there is no reason to think that the subjects were asked about the influence of the relative highness or lowness of the number indicated by the roulette wheel on their judgments. Hence, it seems hasty to conclude that introspection would not reveal such an influence in these cases.

There is one feature of this experiment, however, which Kornblith failed to mention, that might lend more support to the view that introspection would not reveal the anchoring effect to subjects. This feature is that, as Tversky and Kahneman (1982) report, "payoffs for accuracy did not reduce the anchoring effect" (p. 14). It is reasonable to suppose that such potential financial gain would motivate subjects to be more reflective about their judgments than normal. But if careful introspection were to reveal to subjects their susceptibility to the anchoring effect, one would expect subjects aiming at accuracy to attempt to avoid this potential influence on their judgments. However, since payoffs for accuracy did not reduce the anchoring effect, we can conclude that introspection did not uncover its influence. Now even if this argument succeeds, it still provides less support for Kornblith's thesis about introspection than does the previous "position effect"

experiment. This is because, in that latter experiment, subjects were asked explicitly about the influence of relative position on their judgments. However, in the present "anchoring effect" experiment, subjects were not given such a verbal cue as to the nature of this factor which appears to be influential in their judgments. If subjects were to fail to recognize the anchoring effect in their choices even after such a verbal cue, it would be much more reasonable to draw a negative conclusion about the reliability of introspection in such cases. That is, if they could not detect the influence of the anchoring effect on their judgments even after a verbal cue about the effect, then there would be stronger grounds for concluding that introspection does not have access to this factor. Hence, this experiment provides less support for Kornblith's claim than does the previous one.

Kornblith's Case for Premise Two: Experiment #3.

The previous two experiments were intended by Kornblith to show that introspection does not reliably reveal all the factors which influence our judgments. He also appeals to two additional psychological considerations in an attempt to show that introspection may lead us to believe that certain factors have influenced our judgments which have actually had little or no influence on us. The first such consideration is an experiment conducted by Nisbett and Ross (1980, p. 209). Here is Kornblith's description of it:

Nisbett and Ross had subjects read a passage describing a baby's accidental drowning; subjects were then asked to pick out those portions of the passage which they thought were most responsible for its emotional impact. The portions which were indicated by the subjects were just the ones which Nisbett and Ross thought were most responsible for its dramatic effect. When the same passage was given to a different group of subjects with the indicated passages removed, however, they rated it as having the very same emotional impact as the first group (p. 7).

Although Kornblith does not mention this, it seems to me that this experiment is more promising than the previous two as a basis for the claim that introspection of the causes of our mental states is unreliable and misleading. This is because the primary task which the subjects are engaged in is the attempt to introspect what portions of the passage are most responsible for their emotional reaction to the passage. Unlike the other experiments which Kornblith considers, this experiment is designed to test the efficacy of introspection itself. Kornblith infers from the results of this experiment that ". . . introspection simply does not provide us with accurate information about the etiology of our mental states" (p. 7). Although this experiment does indeed provide a stronger basis for this claim than do the previous two, Kornblith is overstepping his bounds by making his conclusion so general. It is more reasonable to make such a claim about cases of the kind described in the passage and similar cases.

Kornblith's Case for Premise Two: Psychological

Consideration #4. The second psychological consideration which Kornblith uses in his attempt to demonstrate that introspection may mislead us concerns the ability of passengers on an airplane to determine what the orientation of the plane is. He does not appeal to any scientific psychological experiments here. It appears that he is merely reporting a personal anecdote. Here is Kornblith's argument:

Passengers on a plane are able to tell when the plane is headed up or down, right or left. This is something, it seems, that one detects visually. The orientation of the plane, however, is not detected by visual means at all, as is revealed by noticing that the very same effect occurs even on night flights where there is no possible visual source of information about the plane's orientation. Information about the orientation of the plane is conveyed not visually, but kinesthetically; one detects the orientation of the plane by detecting the orientation of one's body. Careful introspection is powerless to detect the source of this information. The illusion that the orientation of the plane is detected visually persists even when one knows that it is detected kinesthetically (p. 8).

Before we discuss the merits of this argument, we need to make a correction to Kornblith's psychological claim concerning how people determine the orientation of a plane. Kornblith says that this information is conveyed kinesthetically. However, as Goldman has pointed out to me in conversation, such information is conveyed by means of the functioning of the vestibular organ in the inner ear. According to Chaplin's Dictionary of Psychology (1975), the vestibular sense is a function of the vestibule, which is "a bony cavity in the labyrinth of the inner ear consisting of two

sacs, the utricle and the saccule, which contain hair cells that are affected by acceleration and deceleration of the body and by changes in head position" (p. 563). For ease of exposition, I will couch the discussion of Kornblith's argument contained in the above quotation in terms of the kinesthetic sense. However, it is to be understood that the actual sense responsible for detecting the orientation of a plane one is on is the vestibular sense.

In addition to this factual correction, there are a number of problems with the argument contained in the passage quoted above. In the first place, the fact that one can detect the orientation of a plane even during a night flight when visual cues are absent does not show that ". . . the orientation of the plane (at any time) is not detected by visual means at all." What it shows instead is that visual cues are not necessary for determining the orientation of a plane. It seems reasonable to suppose that visual cues can contribute to the detection of a plane's orientation during a daytime flight. Hence, introspection would be misleading here if it lead one to believe either that visual cues are necessary to detect the orientation of the plane or that visual cues are always entirely responsible (to the exclusion of other kinds of cues such as kinesthetic cues) for one's ability to detect the orientation of a plane.

Given this correction of Kornblith's argument, it is necessary also to recast his claim that ". . . information

about the orientation of the plane is conveyed not visually, but kinesthetically." A more reasonable way of putting this is to say that such information need not be conveyed visually since kinesthetic cues can be sufficient to enable one to acquire this information. What is Kornblith's basis for this claim? Obviously, he cannot legitimately claim to know this on the basis of introspection. Moreover, he has not provided us with experimental evidence for this claim. Perhaps it is based on an inference by elimination: "The only candidate ways to detect the orientation of a plane are visually and kinesthetically. Since it need not be done visually, it must be that kinesthetic cues are sufficient." But it is unclear what the basis is for the first premise of this inference. Hence, I am not sure what psychological evidence warrants Kornblith's claim that information about the orientation of a plane is conveyed kinesthetically.

My own view is that, even if experimental evidence for the claim that kinesthetic cues are sufficient to enable one to detect the orientation of a plane were forthcoming, it would be unnecessary. This is because it seems obvious to me that we can introspect that we can detect the orientation of a plane by detecting the orientation of our bodies. Perhaps what Kornblith has in mind is that while we are looking out a plane's window in daylight it seems to us that our ability to determine the orientation of the plane is due entirely to visual cues. This may be the case. However, it

seems to me that, in the absence of visual cues, such as when we close our eyes, we are able to introspect that kinesthetic cues can enable us to determine the orientation of the plane. If I am right about this, it may be that in his attempt to discredit introspection by means of this example, Kornblith has really isolated a case in which we do have introspective access to some of the sources of our judgments. However, even if I am wrong, it is not clear what psychological evidence warrants Kornblith's conclusion that introspection misleads us about our ability to detect the orientation of a plane. Consequently, this example either supports the opposite conclusion from the one Kornblith defends, or it does not demonstrate anything at all.

Kornblith's Case for Premise Two: Confirmation Bias and Introspection. There is no need to examine the final two psychological experiments which Kornblith summarizes. This is because he does not draw any conclusions about introspection on the basis of the experiments themselves, but instead employs the hypothesis which psychologists have posited to explain the experiments to argue for a thesis about introspection. The experiments in question are studies of hypothesis testing (see Wason and Johnson-Laird 1972, pp. 204-214 and Nisbett and Ross 1980, pp. 170-172). The hypothesis that psychologists believe is confirmed by these experiments is that, in such tasks subjects are prone to exhibit a "confirmation bias". This confirmation bias is

manifested in two distinct ways. Kornblith summarizes these two kinds of confirmation bias as follows:

First, subjects seeking to test an hypothesis tend to seek out confirming instances and do not attempt to check for disconfirming information. Second, when disconfirming information is nevertheless encountered, it tends to be taken far less seriously than confirming information, even when the hypothesis being tested was not antecedently believed (p. 10).

Kornblith uses the confirmation bias hypothesis in a clever argument which purports to show that, not only is introspection unlikely to reveal one's epistemic shortcomings, even to a responsible agent, it also leads a person to develop an inflated and inaccurate opinion of his reliability. Here is the argument: "Most people believe that they are generally reliable in their acquisition of beliefs. Given our natural confirmation bias, this belief about ourselves will tend to be immune to disconfirming evidence. Furthermore, there is an interesting mechanism which supports the operation of confirmation bias with respect to this belief about ourselves: we tend to downplay our mistakes in reasoning as they occur. After all, "everyone makes mistakes" and the discovery of a particular error in reasoning on our part would seem not to call for general doubts about our reliability. Given this lack of emphasis placed on our mistakes, they will tend to be less salient and consequently not as easily recalled as our reasoning successes. Hence, when we reflect on our reasoning ability,

we will tend to remember mostly instances of correct reasoning. Therefore, the more we introspect about the reliability of our belief-forming processes, the more our belief in our reliability is confirmed. That is, the more we try to improve, the more we think we are not in need of improvement. So introspection not only fails to uncover our unreliable processes of belief-formation, it also serves to make us more confident in our reasoning abilities than we ought to be."

This argument has a degree of plausibility, but there is good reason to think that it does not show as much as Kornblith believes it does. The argument may support a claim about the normal use of introspection, but it is arguable that it falls short of demonstrating that the careful, deliberate and responsible use of introspection will fail to reveal some epistemic shortcomings. It seems reasonable to suppose that, if someone was bent on discovering his errors in reasoning, he would attach special salience to any particular mistake which he noticed. Furthermore, he would be quite likely to notice more of his reasoning errors in the first place than would a normally reflective person since he would be consciously on the look-out for them. Hence, it is not unreasonable to think that careful use of introspection could reliably be used to overcome our tendency toward a confirmation bias so that we could attain a more accurate assessment of our general reliability as reasoners. So,

although the normal use of introspection may lead one to an exaggerated opinion of one's general reliability as a reasoner, it is reasonable to think that a careful use of introspection would not.

Moreover, it is not clear that Kornblith has provided enough support for even the weakened form of the "confirmation bias" argument's conclusion concerning only the normal use of introspection. Kornblith's claim that confirmation bias is operative in our normal assessments of our reliability, though somewhat plausible, requires additional experimental support. In particular, what is needed is experimental evidence that confirmation bias is operative in the specific way that Kornblith claims it is in the above argument. Lacking such support, the argument does not decisively establish even the weakened form of its conclusion phrased in terms of the normal use of introspection alone.

Overall Assessment of Kornblith's Case for Premise Two. Now that we have critically reviewed Kornblith's case for premise two of his main argument against the Cartesian thesis that epistemic improvement ought to proceed by means of introspection, what can be said in general about its success? The first three psychological considerations discussed in this section do tend to support Kornblith's claim about the unreliability of introspection vis a vis the relevant applications. However, in each case, the conclusions warranted by the experimental evidence are either weaker or

less general than the one's which Kornblith draws. Recall that Kornblith's overall conclusion on the basis of the psychological evidence he considers is that ". . . our processes of belief-acquisition, and indeed mental processes in general, are largely unavailable to introspection" (p. 14). I think it is clear that, on a more reasonable assessment of the psychological evidence which Kornblith considers, this claim is too strong. A more plausible claim is that not all employments of introspection to determine the sources of our beliefs are reliable. In particular, introspection does not have access to all of the causes of our beliefs and at least sometimes misleads us by making us think something is a cause of one of our beliefs when it is really not. Furthermore, the normal (but not the careful) use of introspection might lead a person to acquire an exaggerated and inaccurate opinion of his general reliability. Stronger claims than this would require substantially more experimental evidence and theoretical support from current cognitive psychology. Given this weakening of the second premise of Kornblith's main argument, the conclusion of that argument must be weakened as well: If one wants to improve epistemically, one must not rely on introspection alone (and perhaps one must not rely chiefly on introspection either). Some other reliable method is required to supplement it. Kornblith concludes his discussion by asserting that this additional method must be cognitive psychology.

The Relevance of Kornblith's Discussion to the Present Project. Now that we have critically examined Kornblith's arguments, we can assess their relevance to our overall program. There are three reasons for including Kornblith's discussion in the present chapter. In the first place, the main structure of Kornblith's argument is quite similar to the main structures of the previous three arguments. Like the other three, Kornblith locates a psychological component in an epistemological thesis which he deems to have a large degree of currency in the epistemological literature. As I have said, the truth of this epistemological position depends upon the truth of the psychological component. Kornblith then appeals to psychological investigation in the attempt to determine the truth of this epistemological thesis. In this case, the psychological evidence which Kornblith uses consists mainly of experiments from current cognitive psychology. He then concludes on the basis of this psychological evidence that the psychological thesis is false and that, consequently, the epistemological thesis which depends on it is false as well. So Kornblith's argument has the same general structure as arguments which are applications of the constraint of psychological realism.

Secondly, if Kornblith's strongest statement of his psychological thesis is true, then it must be concluded that most, if not all, employments of the constraint of psychological realism which appeal to introspective evidence ought

to be ruled out. This is because Kornblith concludes that "our processes of belief-acquisition, and indeed mental processes in general, are largely unavailable to introspection" (p. 14). In particular, this thesis would seem to undermine the applications of the constraint of psychological realism by Cohen and Alston which we surveyed in chapter three. Consider first Alston's argument against Chisholm's view. Recall that his argument relied on reliable introspective access to the way in which our beliefs are formed. More specifically, his use of introspection purported to show that belief-formation is not subject to voluntary control. But if Kornblith's strong claim about introspection is correct, then Alston's psychological premise is subject to serious doubt. What about Cohen's argument against Lehrer's analysis of epistemic justification? Kornblith's psychological claim does not threaten Cohen's use of introspection as much as it does Alston's, for two reasons. First, Cohen places less weight on introspection, using it not to conclude that Lehrer's view is false but instead to shift the burden of proof to Lehrer. This could be accomplished, although with less persuasiveness, even if Cohen's introspective evidence was subject to doubt. Second, Cohen's introspective claim concerns whether or not humans possess certain kinds of mental states. However, Kornblith's disparagement of introspection is limited to its role in determining the properties of our mental processes.

So Kornblith's psychological claim does not directly threaten uses of introspection such as Cohen's. However, if Kornblith's claim about introspection is right, it might seem as if the burden of proof is shifted to someone who would wish to defend any use of introspection whatsoever. At the very least, it would seem that one who appeals to introspection for information about mental states ought to provide cogent reasons for thinking that, although we have little or no reliable introspective access to the nature of our mental processes, we do have such access to the relevant kinds of mental states. Hence, if Kornblith's psychological thesis is correct, some doubt is cast on applications of the constraint of psychological realism like Cohen's as well. Consequently, since Kornblith's psychological thesis constitutes a threat to applications of the constraint of psychological realism which make use of introspection, such as those of Alston and Cohen, it is worth our while to consider his arguments for this claim carefully.

The final reason for including Kornblith's comments on Cartesian epistemology is that his argument is suggestive of possible future applications of the constraint of psychological realism. In particular, it is an illustration of the use of scientific psychological experiments to argue against an epistemological thesis with a psychological component. As for the specific psychological thesis which Kornblith considers, it is not clear that any existing

accounts of epistemic justification require for the possession of justified beliefs that humans be capable of exercising a reliable process of introspection. Of course, if there is such a view, then the success of Kornblith's arguments against the reliability of introspection would demonstrate that such accounts are psychologically unrealistic. There are existing accounts of epistemic justification which require for justified belief that humans be able to introspect, however. Examples of these views are Bonjour (1985), Kornblith (1983) and perhaps Lehrer (1986). Although Kornblith's thesis about the unreliability of introspection may not entail that these views are psychologically unrealistic, it may be that other interesting consequences would follow about them, such as that they are make epistemic justification a relatively weak notion in the sense of not being truth-conducive. So Kornblith's arguments may have implications for some accounts of epistemic justification other than implications about their psychological reality or unreality.

Kornblith's Psychological Thesis and the Use of Introspection in Applying the Constraint of Psychological Realism. What do Kornblith's arguments against the reliability of introspection indicate about the suitability of using introspection as a form of psychological investigation in applying the constraint of psychological realism? If it

could be shown that introspection is completely unreliable with respect to any possible application, then arguments such as those of Alston and Cohen which we examined in chapter three would be inadequate. However, it is clear that Kornblith's arguments fall short of warranting this extremely pessimistic conclusion about introspection. Although his case does demonstrate that a certain Cartesian epistemological thesis is false and in need of modification, it does not show that every use of introspection is unreliable. Moreover, it is quite reasonable to think that introspection is reliable for many purposes and in some cases, even more appropriate than scientific psychological investigation (a point that was argued for in chapter three). For instance, it is plausible to think that the thesis of direct voluntary control of belief is one which one can test introspectively quite easily. All one need do is attempt to will a belief into or out of existence. It is introspectively obvious that this cannot be done. In sum, although Kornblith's arguments support the thesis that introspection is not always reliable, they do not show that every employment of introspection is unreliable. So we need not conclude that introspection ought not to be used for some applications of the constraint of psychological realism. Moreover, reasonable assessment of Kornblith's case for his psychological thesis (premise two) not only fails to show that introspection should be avoided in applying the

constraint of psychological realism, it also supports the claim I made in chapter three that there are some employments of the constraint for which introspection, unsupplemented by cognitive psychology, is not adequate.

Summary. In this section I have critically examined an argument which Kornblith offers against a purportedly Cartesian epistemological thesis in his recent unpublished paper entitled "Introspection and Misdirection." Although Kornblith's argument is not an application of the constraint of psychological realism, it has the same general structure of arguments which incorporate this constraint. Furthermore, it is suggestive of ways in which scientific psychological experiments can be used in applications of the constraint. The epistemological thesis which Kornblith attributes to Descartes and to contemporary epistemologists in the Cartesian tradition is that, in one's search for truth, one ought to begin with introspection. We have seen good reasons for thinking that this thesis is true only if introspection is a reliable method for determining what processes are responsible for causing one's beliefs and whether or not these processes are reliable. Kornblith denies that introspection is reliable for these purposes. He bases this claim on six psychological considerations which include five experiments from cognitive psychology. Although the case against reliable introspection which Kornblith builds on these psychological considerations does demonstrate the need

for modifying the Cartesian thesis, it does not support the strong and general claims that he makes about introspection. A weaker thesis, that introspection is somewhat adequate for the purpose of epistemic improvement, but in need of supplementation by another method, such as cognitive psychology, is more plausible given the psychological evidence which Kornblith adduces. Given this weaker thesis, it is clear that Kornblith's psychological arguments do not imply that applications of the constraint of psychological realism based on introspection are illegitimate. As a matter of fact, the weaker thesis supports our claim that both introspection and scientific psychological methodology are required for the various possible applications of the constraint.

Consequences of Widespread Psychological Unreality

This ends our illustration of two employments of versions of the constraint of psychological realism applied to accounts of epistemic justification, an employment of a version of the constraint applied to theories of rationality, and an argument with a structure similar to those incorporating the constraint. Each of the evaluators claimed that the view he was criticizing is representative of many other similar views in the literature. Along the way, we have seen good reasons for thinking that this is true. It is clear that each of these critics has chosen to

criticize an influential view of a prominent philosopher. Moreover, the evaluators themselves are each highly-regarded epistemologists. I believe we can draw two conclusions about the constraint of psychological realism on the basis of these observations. First, the constraint is already playing an important role in epistemology. Second, if the evaluators' arguments are sound, then leading epistemologists have at one time endorsed influential epistemological views which have nevertheless been psychologically unrealistic. Furthermore, if these arguably psychologically unrealistic accounts are representative of similar views of other epistemologists, and I believe that they are, then it is likely that these other accounts which exist in the literature will fail to satisfy the constraint of psychological realism as well. Consequently, if the arguments of the evaluators are sound, then there is potentially a broad range of psychologically unrealistic epistemological accounts in the literature. Of course, other positions may be found to be psychologically unrealistic on the basis of considerations not advanced by these four critics. The examples provided in this chapter and the last point the way to further applications of the constraint of psychological realism. I would like to end this chapter by considering three consequences which would follow from the widespread psychological realism which is suggested by the success of the four evaluators' arguments against influential and rep-

representative epistemological views.

The Importance of the Constraint of Psychological Realism

What would such a large number of psychologically unrealistic views held by epistemologists indicate about the importance of the constraint of psychological realism and the indispensability of psychological investigation in its application? Assuming that the philosophers who endorse these psychologically unrealistic accounts would also deny J-skepticism(H), there are two possible conclusions which can be drawn about the significance of the constraint. The first conclusion is that there may exist among these philosophers a general disregard of the psychological implications of epistemological theories. This potentially widespread tendency to overlook the psychological components of epistemological views underscores the importance of the constraint of psychological realism as an evaluative tool. Its employment may be necessary to point out some otherwise missed undesirable skeptical consequences of a position. The second possible conclusion is that the acceptance of psychologically unrealistic accounts of epistemic justification by those who reject J-skepticism(H) may be due, not to an insensitivity to the psychological implications of the accounts, but to mistaken beliefs about human psychology. In such cases we can describe the intuitiveness ascribed to these views as conditional intuitiveness, a notion we intro-

duced in chapter three. The intuitiveness involved would depend on a conscious or unconscious adherence to a false psychological thesis. The constraint of psychological realism can play a crucial role in cases of this latter type as well by uncovering these false psychological components and thereby revealing the skeptical consequences of holding the views which contain them.

The Indispensability Thesis

These considerations indicate how the potentially widespread existence of psychologically unrealistic epistemological views emphasizes the important role of the constraint of psychological realism in epistemology. What can be said about the implications for the indispensability of psychological investigation in applying this constraint given the existence of a sizable number of such positions? Whether the potential preponderance of psychologically unrealistic epistemological theories is due to ignorance, error, or both, the broader the extent of such positions, the stronger the defense of the indispensability thesis against Pollock's objection to it discussed in chapter three. This is because a large number of existing psychologically unrealistic views would make it clearer that semantic intuitions are not always a reliable guide to the psychological reality of an account. If intuitions were a reliable guide, fewer people who deny J-skepticism(H) would find accounts intuit-

ive which are actually psychologically unrealistic. Consequently, given the success of the applications of the constraint of psychological realism which we have discussed in this chapter and the representative nature of the accounts undermined, it follows that we have gone a significant way toward demonstrating both the importance of the constraint of psychological realism and the indispensability of psychological investigation in its application.

Although the previous considerations provide good grounds for thinking that the success of our four evaluators, and similar successful applications of the constraint of psychological realism, would be sufficient, when supplemented with the arguments of chapter two, to demonstrate the importance of the constraint of psychological realism and the failure of Pollock's objection to the indispensability thesis, such success would not be necessary to show these things. First, with respect to the importance of the constraint, even if no existing epistemological accounts were psychologically unrealistic, we would still need to apply the constraint to determine this. Accordingly, it is just as important to apply the constraint to determine that an account is psychologically realistic as it is to determine that it is not. Secondly, with respect to defending the indispensability thesis from Pollock's objection, even if no existing epistemological accounts were psychologically unrealistic, it would not follow that intuitiveness is suffi-

cient as a test for psychological reality. The problems with narrow intuitiveness, such as intuitional variability, would still make it unreliable as an indication of empirical theses concerning the psychological reality of an epistemological position. So the success of the evaluators' arguments and similar arguments would help considerably to show the importance of the constraint of psychological realism and to defend the indispensability thesis from Pollock's objection, but such success would not be necessary.

Motivation for My Proposal

I will mention just one more consequence of the success of the evaluators' employments of the constraint of psychological realism. If the critics' arguments are sound, then they can be viewed as providing some motivation for the approach toward an account of epistemic justification which I will sketch in chapter five. This motivation would by no means be an argument by elimination. The epistemological views examined by these evaluators do not jointly exhaust logical space. Even if the downfall of these positions results in the defeat of other related accounts, it is not clear that all possible plausible alternatives to my proposal will have been eliminated. However, such a result would constitute some clearing of the ground, and together with additional considerations which I will adduce in chapter five, this result would contribute to the case for my

proposal. I hope to show that an account of the sort that I will sketch would satisfy both the constraint of intuitiveness and the constraint of psychological realism. I should reiterate, however, that the main purpose of this chapter is not to demonstrate that four epistemological views are psychologically unrealistic, but to illustrate the application of the constraint of psychological realism. This is why I have not been concerned to determine if the four critics have accurately represented the views which they argue against. Any motivation for my own proposal would be a welcome, but secondary result. I will discuss the precise nature of such a potential motivation in the next chapter.

Summary

In this section, I have argued that, assuming that the four employments of the constraint of psychological realism which we have examined are successful, a stronger case is provided for two theses for which I argued in chapter two. First, the importance of the constraint of psychological realism as an evaluative tool in epistemology is underscored. Second, psychological investigation is seen to be a more reliable method than appeal to semantic intuitions for determining whether an epistemological account is psychologically unrealistic. The case for these two claims is especially strong in light of the representative nature of the four views which would be undermined. Of course, the

complete failure of or defect in any or all of the four critical arguments would not necessarily defeat these two theses. It is likely that other applications of the constraint of psychological realism would succeed. Furthermore, even if there were no psychologically unrealistic epistemological theories, we would need to apply the constraint to determine this. Future proposals would need to be tested as well. A third consequence of the soundness of the critics' arguments in this and the previous chapter would be some motivation for a proposal I will make in chapter five.

CHAPTER FIVE
TOWARD A PSYCHOLOGICALLY REALISTIC ACCOUNT OF
EPISTEMIC JUSTIFICATION

In chapters three and four, we saw how the constraint of psychological realism has been and can be used as a critical tool for testing candidate accounts of epistemic justification. In this chapter, I will illustrate a way in which the constraint can be employed to guide research into the nature of epistemic justification. The conclusions of the previous three chapters underscore the need for this research to be guided by considerations of psychological reality. In chapter two, I argued that no psychologically unrealistic account of epistemic justification is true. An account of justification is psychologically unrealistic just in case it requires for the possession of justified beliefs some psychological capacity which human beings do not possess. I also argued that this constraint is important since it is quite reasonable to think, on intuitive grounds, that the correct account of epistemic justification has an interesting psychological component. Any candidate account which incorporates such a psychological component would need to be tested for psychological reality. I also argued for what I called the indispensability thesis. This is the view that psychological investigation is required for testing accounts

of epistemic justification for psychological reality. As I mentioned above, the previous two chapters illustrated a critical use of the constraint of psychological realism. They also provide more support for my contentions that the constraint is important and that psychological investigation is indispensable in its application. In addition, they suggest a tendency on the part of some philosophers to neglect the results of psychological investigation which have a bearing on their epistemological views or to hold arguably mistaken psychological views on which the intuitiveness of their epistemological views depends. Given all of these conclusions, it is quite reasonable to look for a research program concerning the nature of epistemic justification according to which one can be guided by considerations of psychological reality in the construction of an account of justification as well as in testing a constructed account.

Cherniak's argument for a psychologically realistic thesis about minimal rationality, which we examined in chapter four, provides a useful model of guidance by considerations of psychological reality in epistemological theory construction. Recall that Cherniak argued that, given the short term/long term memory distinction, the long term memory must be compartmentalized in some fashion. Cherniak's strategy here is to start with a general psychological thesis, which has a substantial degree of support, and to engage in thought experiments about further more specific

ways in which the human belief system must be designed in order for a human being to be minimally rational. I argued that this reliance on intuitions is not sufficient for determining that the thesis is correct. However, as I urged in chapter one, satisfaction of the constraint of intuitiveness is necessary for determining that an epistemological account is correct. We are now in a position to see the way in which these two constraints can be applied in tandem in order to yield an account which is both intuitive and psychologically realistic. In the beginning stages of research, we take what we know about the mind and do thought experiments which will generate intuitive psychological components of an account of epistemic justification. At the same time, we continue to consult the results of psychological investigation to test the intuitively-generated theses. Discovery of psychological unreality will call for theoretical modification. Confirmation of psychological reality will of course constitute satisfaction of the constraint of psychological realism. In this way, psychological considerations can be used alternately to guide the construction of an account of epistemic justification and to test it for psychological reality at each successively more determinate stage of theory construction.

An Account-Schema Which Satisfies Both Constraints

What general psychological thesis or theses can

serve as a starting point in the construction of an account of epistemic justification? What we need is a psychological component or psychological components which will constitute a schema for the correct account of justification. Such a schema must be intuitive and psychologically realistic, since these are two constraints which we have argued must be satisfied by the correct account of epistemic justification. In addition, the psychological component(s) must be general enough to be species-neutral. That is, they must be general enough to encompass any possible psychological species which is capable of possessing justified beliefs. Here intuition must be consulted, but the deliverances of intuition are always subject to challenge by considerations of psychological realism. For instance, intuition may suggest that possession of psychological capacity C is required by every psychological species for justified beliefs. However, if psychological investigation reveals that humans do not possess C, then the intuitionally-derived thesis will require modification. In such a case, we must look for another general thesis which is also both intuitive and psychologically realistic.

I believe that we can locate a relatively uncontroversial, general, intuitive and psychologically realistic set of psychological components for an account of epistemic justification by focusing on the maintenance of the belief system. Over a period of time, a person's belief system is in

constant flux. New beliefs are constantly being added, many existing beliefs are sustained, for shorter or longer periods of time, and many beliefs are eliminated from the system. A change in belief status in any of these three ways occurs by means of the operation of some appropriate cognitive process. In chapter three, I labeled such processes "belief processes." The three kinds of belief processes which result in these three changes in belief status can be called processes of belief acquisition, belief maintenance, and belief elimination, respectively. I intend for these cognitive processes to be construed quite broadly to include any mental operations which contribute to yielding the belief status in question. In addition, although the working of these processes directly affects only beliefs, their inputs can include mental states other than beliefs as well as beliefs. I will argue that these psychological components can be incorporated into a schema for accounts of epistemic justification which is intuitive and psychologically realistic.

Our goal is to come up with a schema for accounts of epistemic justification which would serve, as much as possible, to make satisfaction of the constraint of psychological realism non-accidental by accounts of justification which are instances of the schema. More generally, we need a schema which will help direct us to an account which satisfies both meta-theoretical constraints which we have

endorsed. In light of the above remarks about belief system maintenance, we might be inclined to consider the following schema:

- (5.1) For any person S, proposition p and time t,
S's belief that p at t is epistemically
 justified if and only if S's belief that p at
t is held as the result of the actual maintenance of S's belief system.

Now if belief system maintenance involves the interaction of the three belief processes mentioned above, it is quite reasonable to hold that (5.1) is psychologically realistic. Both introspection and current cognitive psychology provide ample support for the psychological thesis that maintenance of the human belief system involves processes of belief acquisition, maintenance and elimination. However, even though (5.1) is arguably psychologically realistic, it obviously does not satisfy the intuitiveness constraint. This is because there are numerous counterexamples to the sufficient condition for epistemically justified beliefs entailed by (5.1). On many occasions, the actual functioning of the human belief system yields epistemically unjustified beliefs. For instance, the belief acquisition process which is commonly called "wishful thinking" yields epistemically unjustified beliefs. Consequently, although the epistemic justification of a belief may have something to do with maintenance of the belief system, it does not involve

just any actual process having to do with belief maintenance.

It is clear that we have some unjustified beliefs. It is equally clear that we have some justified beliefs. Thus, some interactions of belief processes result in justified beliefs, and some result in unjustified beliefs. This suggests the following schema for accounts of epistemic justification which incorporates belief system maintenance capacity as a psychological component:

- (5.2) For any person S, proposition p and time t,
S's belief that p at t is epistemically
 justified if and only if S's belief that p at
t is held as the result of the correct
 maintenance of S's belief system.

Substitution of the word 'actual' with the word 'correct' results in a schema which, at this level of abstraction, is not only psychologically realistic, but avoids the intuitive difficulties of (5.1) mentioned above as well. The word 'correct' captures the idea that only belief processes and interactions of belief processes which are of the right kind will confer justification on a belief.

I intend for (5.2) to be relatively uncontroversial, at least among philosophers who believe that the correct account of epistemic justification has an interesting psychological component. For instance, as I understand it, it is consistent with both types of accounts of epistemic jus-

tification to which the constraint of psychological realism was applied in chapters three and four, and with other types of views as well. In the first place, on a responsibility view, correct maintenance of one's belief system would involve either direct or indirect voluntary control of belief acquisition, maintenance and elimination. Of course, as we have seen, the first type of voluntary control is arguably not within the psychological repertoire of human beings. Secondly, on a positive holistic internalist coherence view, such as that of Bonjour (1985), correct maintenance of one's belief system involves both the maintenance of a sufficient degree of coherence among one's beliefs, by means of a certain degree of (probably indirect) voluntary control of one's belief processes and a "reflective grasp" of the fact that one's entire belief system is coherent to this degree. Introspection, on this view, seems to serve the function of monitoring the maintenance of the belief system. I think that the discussions in chapters three and four cast Bonjour's theory, and others like it, such as Lehrer's in (1974), into doubt because of their arguable failure to satisfy the constraint of psychological realism. Consequently, I will not consider any such accounts as instances of (5.2) in what follows. Nor will we consider responsibility views which require the direct voluntary control of beliefs for the possession of justified beliefs. These exclusions constitute some narrowing

of alternatives for the correct account of epistemic justification. Hence, they comprise a degree of motivation for a proposal I will make toward the end of this chapter. (5.2) is also consistent with reliabilist accounts of epistemic justification. Correct belief-system maintenance on a reliabilist view includes the operation of belief processes and interactions of belief processes which tend to yield mostly true beliefs.

I have argued that (5.2) is psychologically realistic since there is ample psychological evidence that our belief systems are maintained in the three-fold manner described above. In addition, it avoids the intuitive difficulty I raised against (5.1). However, a person who is inclined to doubt that the correct account of epistemic justification has a psychological component beyond the component which requires the capacity for beliefs for the possession of justified beliefs might want more intuitive motivation for the stronger psychological component contained in (5.2). In order to provide this intuitive motivation, it will be necessary to make (5.2) somewhat more determinate. After I have motivated a more determinate form of (5.2), I will consider a few alternative instances of (5.2) as candidate accounts of psychologically realistic, intuitive accounts of epistemic justification.

The following schema displays more clearly the kind of interaction between the three belief processes which is

necessary and sufficient for the correct maintenance of a person's belief system:

(5.3) S's belief that p at t is held as the result of the correct maintenance of S's belief system if and only if:

(a) S's belief that p at t is the output of a belief acquisition process or belief maintenance process of the right kind which is among S's cognitive capacities, and

(b) It is not the case that S's belief that p at t would have been eliminated by a belief elimination process of the right kind which is possessed by S.

This schema is quite similar to the general structure of Goldman's accounts of epistemic justification in (1979) and (1986). We will see presently how his reliabilist views can be construed as instances of it.

Clause (a) of (5.3) might be called the "causal component" of an account of epistemic justification. A person's beliefs are either causally generated or causally sustained by means of the two belief processes to which (a) refers. Psychological processes which either produce or sustain beliefs can be thought of as causal mechanisms taking reasons or grounds as inputs and yielding beliefs as output. Some of these processes take beliefs (reasons) as

input. Others take non-belief mental states (grounds) as input such as perceptual states. Since perceptual states originate in the external world, some of these belief-causing processes can be considered to be externally-linked, in the sense of being connected with the external world. In addition, clause (a) makes it clear that (5.3) is a schema for accounts of actual epistemic justification and not counterfactual epistemic justification, as these notions were defined in chapter one. The psychological processes to which clause (a) refers provide a causal link between S's belief that p and the justifying reasons or grounds for this belief. This goes some way toward an understanding of the basing relation between a justified belief and its justifying reasons or grounds.

Clause (b) of (5.3) constitutes a kind of coherence component of our schema. Belief elimination processes take some of a person's beliefs as input and yield as output the elimination of some of the person's beliefs. Correct maintenance of a belief system requires a certain amount of coherence among a person's beliefs. Beliefs which do not cohere in the right sort of way are eliminated by belief elimination processes of the right kind in a correct system of belief maintenance. Belief elimination processes are not externally-linked as are some processes of belief acquisition and maintenance. Their function is not to maintain connections between the cognitive system and the world but

to establish and maintain the right sorts of connections between the mental states in the cognitive system. Cherniak's argument in chapter four against a Quineian model of the belief system makes it doubtful that this coherence function is holistic. Consequently, considerations of psychological realism will guide the theorist away from such a holistic model when she is spelling out the details of what comprises the right kind(s) of belief elimination processes. Much work remains to be done in this area. This is a research project which would benefit greatly from psychological investigation used both to guide and to test the construction of a completely determinate account of epistemic justification.

This coherence component captures the intuition that epistemic justification is defeasible. I introduced the notion of defeasibility briefly in chapter one in the context of discussing the structure of epistemic principles. According to the conjunction of (5.2) and (5.3), although a belief may be produced or sustained by the right kind of belief process, it may fail to be epistemically justified because a correct belief elimination process would have eliminated the belief from the system. There are two general kinds of belief elimination processes. One of them takes as input beliefs (reasons) for thinking that a belief is false. The other takes as input beliefs (reasons) for thinking that the reasons or grounds on which the belief in question is

based are not good reasons or grounds. Another way of describing this latter type of belief elimination process, although accurate, does not necessarily construe the output of the process as it would be represented in the mind of the person. This is to say that the process takes as input beliefs (reasons) for thinking that the belief process responsible for causally generating or sustaining a belief is a process which is not of the right kind. This sketch of clauses (a) and (b) of (5.3) is sufficient for present purposes. It is now time to provide intuitive motivation for the necessity of each of these clauses for the epistemic justification of belief.

Consider first clause (a). Why isn't it sufficient for epistemic justification merely to have undefeated grounds or reasons for a belief without there being a causal connection between the two? An example due to Gilbert Harman makes it clear that a causal connection is necessary:

Albert believes that he will not pass his ethics course. He has excellent reasons for believing this, because he failed the midterm examination, he has not been able to understand the lectures for several weeks, and the instructor is known to fail a high percentage of his students. Despite the fact that, like many other students, Albert does not appreciate the force of such reasons, he is influenced by something else. During a class discussion early in the term, he gave an emotional speech in favor of existentialism; and, although his instructor said nothing at the time, he seems to advocate linguistic analysis. Albert reasons that, since analytic philosophers hate existentialists, the instructor will remember his remarks, hold them against him, and fail him in the course (1973, p. 24).

It is clear from this example that, even though Albert believes that he will fail his ethics course and has excellent and presumably undefeated reasons for believing this, he is not justified in believing that he will fail. This is because, even though Albert possesses the right reasons for this belief, the belief is caused by the wrong reasons. In order for a belief to be justified, it must be caused by the right reasons.

Given this intuitive motivation for clause (a) of (5.3), why is it also necessary that a belief be undefeated? That is, why is clause (b) of (5.3) necessary for epistemic justification? One's belief is defeated when either one has reason to believe that it is false or one has reason to believe that the process which produced the belief is not of the right kind. An example due to Goldman makes it intuitive that a belief is justified only if it is undefeated:

. . . Millicent in fact possesses her normal visual powers, but has cogent reasons to believe these powers are temporarily deranged. She is a subject of a neurosurgeon's experiments, and the surgeon falsely tells her that current implantations are causing malfunction in her visual cortex. She is persuaded that her present visual appearances are no guide at all to reality. Yet despite this belief, she continues to place credence in her visual percepts. She ignores her well-justified belief in the incapacitation of her visual faculty; she persists in believing, on the basis of visual appearances, that a chair is before her, that the neurosurgeon is wearing a yellow smock, and so on (1986, pp. 53-54).

In this example it is intuitively clear that Millicent is not justified in believing that a chair is before her or

that the neurosurgeon is wearing a yellow smock, even though these beliefs were produced by her normal visual powers. This is because she has reason to believe that her visual processes are malfunctioning. On the reasonable assumption that normal vision is a psychological process of the right kind, it follows that clause (a) of (5.3) is not sufficient for justified belief. Millicent's perceptual beliefs are not justified because she has reason to believe that they are not being produced by the right kind of psychological processes. In the absence of this defeating belief, she would be justified. Hence, clause (b) of (5.3) is necessary for epistemic justification.

Now that I have provided intuitive motivation for clauses (a) and (b) of (5.3), the conjunction of (5.2) and (5.3) is arguably an intuitive and psychologically realistic schema for accounts of epistemic justification. Epistemic justification is a function of the correct functioning of a person's belief system. Correct functioning is cashed out in terms of the appropriate interaction between the processes of belief acquisition, belief maintenance and belief elimination. Moreover, these three processes must be the right kind of processes of belief acquisition, maintenance and elimination. We are left with the following question: "What property, if any, distinguishes belief processes of the right kind from those which are not of the right kind, for purposes of epistemic justification?" In other words,

what is it that distinguishes justification-conferring belief processes and belief process interactions from those which do not confer justification? I will consider some answers to this question which can be found in current epistemological literature. An answer to this question would constitute a further step toward a fully determinate account of epistemic justification. At each such step, psychological investigation should be consulted both to guide intuition and to test for psychological reality.

Reliabilist Accounts of Epistemic Justification

The first candidate answer to the question concerning what distinguishes processes and process interactions which confer epistemic justification on beliefs from those which do not is Reliabilism. I will consider first the form of this very influential response which is proposed and defended by Goldman in his seminal article "What is Justified Belief?" (1979). Goldman's account of epistemic justification in (1986) incorporates important variations on his theory in (1979). However, for present purposes, the earlier form of his view will suffice. I will introduce some of the elements of his theory in (1986) later. Goldman does not employ exactly the same terminology contained in (5.2) and (5.3) to state his view. However, phrasing his position in these terms will not alter its essential features greatly. Goldman's view can be characterized as the con-

junction of (5.2), (5.3) and the following:

- (5.4) B is a belief process of the right kind if and only if B tends to produce a sufficiently high ratio of true beliefs.

I will call the view which consists of the conjunction of (5.2), (5.3) and (5.4) "classical reliabilism."

Given the two meta-theoretical constraints which we have endorsed, we must ask two questions about classical reliabilism and other reliabilist views. These are: (1) Is any reliabilist account psychologically realistic? and (2) Does any reliabilist account satisfy the intuitiveness constraint? Classical reliabilism will be the first reliabilist view we will consider in answering question (1).

Is Reliabilism Psychologically Realistic?

Let us consider the question concerning psychological reality first. Hilary Kornblith, in (1985a, pp. 11-12), gives an example of a way of testing classical reliabilism. He calls the method of assessment used in the example an "empirical test." This empirical test is akin to the constraint of psychological realism, although Kornblith leaves open the question concerning whether J-skepticism(H) is true. As I have defined it, the constraint of psychological realism is the thesis that any account of epistemic justification is false which entails J-skepticism(H) because of requiring a psychological capacity for the possession of

justified beliefs which humans do not have. My argument for this constraint is based on a denial of J-skepticism(H). Even though Kornblith's empirical test does not require this rejection of J-skepticism(H) and is consequently not equivalent to the constraint of psychological realism, it is nevertheless similar to the constraint because it involves determining whether an account of epistemic justification is psychologically realistic by means of an appeal to psychological investigation. However, Kornblith stops short of endorsing the indispensability thesis. He construes his empirical test as at least (and perhaps only) a more expedient way of assessing an account of epistemic justification such as classical reliabilism for psychological reality than a test which consists exclusively of a priori considerations. Given these qualifications, let us consider how Kornblith subjects classical reliabilism to an empirical test for psychological reality.

An expanded and embellished version of Kornblith's example is as follows. Given an endorsement of Goldman's analysis of reliabilism as stated in (5.2) - (5.4), one may find oneself reasoning in the following way:

(5.5) Induction is a reliable belief-acquisition and belief-maintenance process (implicit or explicit psychological assumption).

(5.6) Jones believes that p at t on the basis of induction and there is no reliable process of

belief elimination possessed by her which
would result in her not believing that p at t.

- (5.7) Therefore, Jones is epistemically justified in believing that p at t (this follows from the conjunction of classical reliabilism with (5.5) and (5.6)).

Furthermore,

- (5.8) It is intuitive that Jones is epistemically justified in believing that p at t.

Suppose that one rehearses this line of reasoning for many of Jones's inductively-produced beliefs and never encounters an unintuitive result. However, suppose that, in addition, the best of current cognitive psychology yields an overwhelming amount of evidence for the view that many, if not all, of our inductive processes are unreliable. It provides evidence, in other words, for the falsity of (5.5).

This pessimistic psychological thesis is not entirely implausible. The work of Kahneman and Tversky (see Kahneman, Slovic and Tversky 1982) and others such as Nisbett and Ross (1980) suggests that something like this may be the case. Assuming that you have encountered this psychological evidence, you have three alternatives at your disposal. First, you could continue to endorse Goldman's analysis of epistemic justification and accept the psychological evidence which indicates that inductive processes are unreliable. This move would commit you to skepticism, at least

with respect to the human capacity to possess justified beliefs which are inductively-produced. Second, you could find a flaw in the psychological evidence concerning induction. Third, you could reject or revise Goldman's account of justified belief. These latter two options would enable you to avoid the skeptical conclusion.

The first thing to notice about this example is that one could go with the first, skeptical alternative without necessarily running afoul of the constraint of psychological realism as I defined it in chapter two. This is because the kind of skepticism involved in the above example differs from the kind of skepticism the denial of which generated the constraint of psychological realism in chapter two. The skepticism in the example is skepticism with respect to the justification of inductively-produced beliefs. It leaves open the possibility of there being justified beliefs which are produced by cognitive processes other than inductive processes. The kind of skepticism the rejection of which leads to the constraint of psychological realism is a skepticism about the epistemic justification of any beliefs possessed by humans. Goldman's distinction in (1986, p. 29), concerning the scope of skepticisms, is helpful here. He distinguishes between "local skepticisms," which deny only the justification of some kinds of beliefs, and "global skepticisms," which deny the justification of all beliefs. Skepticism about the justification of inductively-produced

beliefs is a local skepticism. On the other hand, the kind of skepticism disallowed by the constraint of psychological realism as defined in chapter two is global skepticism with respect to the epistemic justification of beliefs. The former, local skepticism denies only that humans are capable of having justified inductive beliefs. The latter, global skepticism denies the human capability of having any justified beliefs whatsoever, regardless of the kind of process which produces the beliefs. The first three epistemological views scrutinized in chapters three and four could be considered to entail this kind of global skepticism.

Given this, it is clear that, in order to determine whether classical reliabilism satisfies the constraint of psychological realism, as defined in chapter two, one would need to examine the reliability of every kind of cognitive process involved in the production, maintenance and elimination of beliefs. Classical reliabilism would fail to satisfy the constraint only if it could be determined that no human belief process of any kind is sufficiently reliable to qualify as being justification-conferring according to the view. Of course, additional constraints similar to the constraint of psychological realism could be generated on the basis of denying particular local skepticisms. The plausibility of such constraints would depend on our convictions about the existence of justified beliefs in the

various domains in question. If our conviction is quite strong that, for instance, humans have justified perceptual beliefs, then a constraint could be generated according to which any account of epistemic justification is false which entails skepticism about the human capacity to possess justified perceptual beliefs.

Let us call the version of the constraint of psychological realism which I proposed and defended in chapter two a "global" constraint, since it holds that no account of epistemic justification is correct which entails global skepticism about the justification of human beliefs. In addition, we can call versions of the constraint of psychological realism which deny only local skepticisms "local" constraints. Classical reliabilism can be tested for psychological reality by means of the global constraint endorsed in chapter two or by means of whatever local constraints are found to be acceptable. Now before one can administer either sort of test, it is necessary, not only to find psychological evidence concerning the reliability of various kinds of belief processes, but also to determine, on non-psychological grounds, how reliable a process needs to be in order to qualify as a justification-conferring process. So far, our statement of classical reliabilism as the conjunction of (5.2), (5.3) and (5.4) leaves this latter question unanswered. At this point, our formulation of classical reliabilism is still only an account-schema.

Consequently, only more determinate formulations of reliabilism can be tested for psychological reality. Once such a determinate account is proposed, we can turn to psychological investigation to find out if human belief processes are sufficiently reliable, given the precise specification of the required truth ratio incorporated in the account. If no belief processes are sufficiently reliable, then the account is psychologically unrealistic. If some belief processes are sufficiently reliable, then the account is psychologically realistic.

Determinate formulations of classical reliabilism can be tested for psychological reality on the basis of either the global constraint of psychological realism or on the basis of various local constraints. We are now in a position to consider the implications of the (global) constraint of psychological realism for non-classical versions of reliabilism. Goldman, in (1986), considers a group of "truth-ratio" accounts of epistemic justification, some of which fall into the reliabilist category. In order to consider the prospects of these accounts for satisfying the constraint of psychological realism, we need to look at the general schema of which each of them is an instance. Goldman states his view of justification in (1986) schematically as follows:

- (5.9) $\underline{S}$'s believing $\underline{p}$ at $\underline{t}$ is justified if and only if:

(a) S's believing p at t is permitted by a right system of justificational rules (J-rules),

and

(b) this permission is not undermined by S's cognitive state at t.

For present purposes, we can provide the following analysis-schema for a right system of J-rules (it is quite similar to, but somewhat more general than the one Goldman offers in (1986, p. 106)):

(5.10) A J-rule system R is right if and only if R permits certain (basic) psychological processes, and the instantiation of these processes would result in a truth ratio of beliefs that meets some specified threshold which is fixed either absolutely or relative to humanly available processes.

(5.10) is a schema which can be instantiated by all of the "criteria of J-rule rightness" which Goldman endorses.

Recall from chapter one that J-rules are epistemic principles. However, unlike the example epistemic principle given there, J-rules do not permit beliefs from an epistemic standpoint. Instead, they permit what I have called "belief processes" from an epistemic standpoint. Given this, we could rephrase (5.9) as the conjunction of (5.2), (5.3) and the following:

(5.11) B is a belief process of the right kind if and only if B is permitted by a right system of J-rules.

I believe this latter formulation captures roughly what Goldman intends by (5.9).

As it is stated, (5.10) can be instantiated by two general kinds of criteria of J-rule rightness. These two kinds of criteria are ones that fix the truth ratio relative to human psychological capacities and ones which fix the truth ratio without regard to these capacities (these are the "absolute" criteria). The former kind of criteria are called "resource-relative" criteria, and the latter kind are called "resource-independent" criteria. According to Goldman,

A resource-independent criterion fixes an acceptable truth ratio without regard to the resources of the (type of) cognitive system in question. A resource-relative criterion fixes an acceptable truth-ratio as a function of the target cognitive system's resources, in the present case, human cognitive resources (1986, p. 104).

Furthermore,

Resource-relative criteria are likely to be comparative. One such criterion is a maximizing criterion. It would say that a rule system is right if and only if the processes it permits would maximize truth ratio, that is, maximize relative to humanly available processes. This is what some writers call a criterion of "doing the best one can." A different sort of resource-relative criterion is a satisficing criterion. For example, a criterion could say that a rule system is right just in case the processes it permits have a truth ratio of a level humans commonly attain, or that it is easy for them to attain (1986, p. 104).

Are truth-ratio accounts of epistemic justification which incorporate a resource-relative criterion for fixing the required truth-ratio psychologically realistic? Yes. The reason for this is that, given the psychological reality of our schema for accounts of epistemic justification the conjunction of (5.2) and (5.3)), if an instantiation of this schema is resource-relative, then it is psychologically realistic. This is easy to demonstrate. In chapter two, I defined an account of epistemic justification as being psychologically realistic if and only if it does not require for the possession of epistemically justified beliefs some psychological capacity which humans do not have. Given Goldman's remarks in the above two quotations, it seems reasonable to conclude that a criterion of J-rule system rightness (and the account of epistemic justification which incorporates it) is resource-relative only if what it requires for the possession of justified beliefs is within the capacity of humans to attain. Consequently, if an account of epistemic justification is resource-relative, then it is psychologically realistic. So, assuming that the schema which comprises the conjunction of (5.2) and (5.3) is psychologically realistic, any resource-relative instantiation of this schema will be psychologically realistic as well. Of course, psychological investigation will be required for determining whether any particular alleged resource-relative account is psychologically realistic. For

instance, in the case of truth-ratio accounts, psychological investigation is required for determining what the truth-ratios are of the outputs of the various human belief processes either in terms of the capacities of these processes (for a maximizing criterion) or in terms of the typical or normal operation of these processes (for a satisficing criterion).

Resource-relative accounts of epistemic justification are psychologically realistic. Are any resource-independent truth-ratio accounts of epistemic justification psychologically realistic? A resource-independent account of epistemic justification is psychologically realistic if and only if some human belief processes have the capacity to produce the truth-ratio of beliefs which is required for justification by the account. Given a precise specification of the required truth-ratio, one would need to consult psychological investigation to determine if there is such a process. Psychological investigation will reveal which candidate resource-independent accounts of epistemic justification are psychologically realistic and which are not. Those accounts which fix the truth-ratio higher than the capacity of any human belief process to attain are psychologically unrealistic.

Now if it should turn out that human belief processes are generally unreliable (a state of affairs which is unlikely, especially given Goldman's remarks throughout the

second half of (1986)), then there will be no reliabilist truth-ratio views which are psychologically realistic, since such views would depend on human cognitive processes being at least minimally reliable (that is, having a truth-ratio of at least .5; Goldman makes this point on p. 105 of 1986). Instead, there would be only non-reliabilist truth-ratio views. I will assume that Goldman's careful examination of various cognitive processes for reliability in the second half of (1986) demonstrates that some reliabilist view is psychologically realistic (see for instance pp. 189-191, 210-214, 219, 246, 320-323). This is because at no point does he encounter evidence which indicates the general unreliability of a human cognitive process. Furthermore, he adduces adequate evidence which shows that many human cognitive processes are at least minimally reliable. It follows from this and what I said above that some reliabilist views are psychologically realistic. This answers our first question about reliabilism. The second question concerns whether a form of reliabilism satisfies the constraint of intuitiveness. I will point out, in this section and the next, why some people have thought that it does not and why others have insisted that it does.

Does Reliabilism Satisfy the Intuitiveness Constraint?

Stewart Cohen advances a counterexample to reliabilism which purports to show that production or maintenance by

a reliable process is not necessary for justified belief (1984, pp. 281-284). Borrowing the scenario from Descartes, Cohen asks us to consider a possible world in which a malevolent demon arranges things in such a way that our cognitive processes (e.g. perception, memory and inference) are unreliable. On a reliabilist view, of course, the inhabitants of such a world would have no justified beliefs. The demon also sees to it that the phenomenal experiences of the demon-world dwellers are coherent and normal. That is, they are just the sort of perceptual experiences one would expect them to have if all of their beliefs were in fact true. Consequently, these people have no reason to think that their cognitive processes are unreliable. In opposition to reliabilism, Cohen insists that the demon-world denizens can have justified beliefs, in spite of their production by unreliable processes. His argument amounts to the claim that reliabilism fails to satisfy the intuitiveness constraint.

Cohen introduces two cognizers, A and B, who live in the demon world, to make his case. A is a good reasoner, and B draws all of his inferences on the basis of confused reasoning, wishful thinking, reliance on emotional attachments, guesswork, etc. Cohen argues that we surely want to conclude that, in spite of the fact that no perceptual process is reliable in such a world, A is epistemically justified in holding his beliefs, whereas B is not. This is

because we would want to say that A has the right to believe what he does, whereas B does not. A has been epistemically responsible with the evidence at his disposal, whereas B has not. The actual objective truth-ratio of their beliefs should not affect these conclusions. Consequently, production or maintenance by a reliable cognitive process is not necessary for justified belief.

In (1986, pp. 106-113), Goldman attempted to avoid this counterexample by proposing a reliabilist theory which does not make justifiedness a function of the reliability of cognitive processes relative to whatever world the person possessing those processes happens to be in (in Cohen's example this is a demon world). Instead, he makes justifiedness a function of reliability in normal worlds. Normal worlds, according to Goldman, are worlds which are consistent with our general beliefs about the actual world. Although this proposal does indeed avoid Cohen's objection, it has serious problems which Goldman himself recognizes in (unpublished, pp. 15-16). Chief among these problems is the difficulty posed by a lack of constraints on general beliefs. Which general beliefs about the actual world are relevant? Do these beliefs have to be justified themselves? If not, what is the epistemic value of reliability relative to them (see Pollock 1986, p. 117)? Peoples' general beliefs about the world vary, perhaps markedly. We cannot appeal to this entire set of (perhaps inconsistent) general

beliefs to specify normal worlds, so how do we go about choosing which ones to use? Another problem is noted by Pollock:

. . . if our general beliefs are not formed by processes reliable in the actual world, then there is no reason to regard beliefs formed by processes reliable relative to those general beliefs as being probably true. Thus this strategy for avoiding (counterexamples such as Cohen's) appears to forsake the general intuitions that made reliabilism attractive in the first place (1986, p. 117).

It is for some of these reasons and others that Goldman, in (unpublished) abandons the "normal worlds" approach and opts for a qualified form of the view that the justifiedness of a belief in a given world is a function of the reliability of the processes which produce the belief in that world. This position appears to run afoul of Cohen's counterexample. Goldman, in (unpublished), attempts to handle the intuitions which drive Cohen's example by arguing that what distinguishes reasoners A and B in the demon world is that, although neither of them is justified in the full-blooded reliabilist sense (neither is "strongly justified" in Goldman's terminology), A is justified in a non-reliabilist sense and B is not so-justified (in Goldman's terminology, A is "weakly justified" and B is not). This latter notion of weak justifiedness is characterized by Goldman as "mere blamelessness." That is, where "ill-formed" in this context means "unreliable," it is a notion of "ill-formed-but-blameless belief" (p. 8).

Given this response of Goldman, the crucial question concerns whether strong justification, which is characterized by his reliabilist view, really is a notion of epistemic justification. Here intuitions diverge. Some may be inclined to say that what Goldman calls "strong justification" is really not epistemic justification (or even a type of epistemic justification, assuming that there is more than one type), but perhaps instead a separate component of knowledge. This position might be taken by someone who believes that Cohen's counterexample ought to be avoided by the correct account of epistemic justification. On the other hand, some may not be satisfied with Goldman's notion of "weak justification," which is a notion of mere blamelessness, or similar accounts which satisfy the intuitions which drive Cohen's counterexample, as an account of epistemic justification. This notion might seem too weak to capture the notion of epistemic justification. Is strong justification too strong for epistemic justification? How one answers this question will depend on whether one thinks that the correct account of epistemic justification must avoid Cohen's counterexample. I will not take a position on this issue. Instead, I will investigate alternatives to Goldman's proposal in (unpublished) which fall into the category I will call "design accounts of epistemic justification." One of them is a version of reliabilism. These accounts are promising in terms of satisfying the constraint

of psychological realism. It also seems reasonable to expect that some such account will satisfy the intuitiveness constraint. Of course, whether this is the case depends on one's intuitions concerning, among other things, the epistemic status of beliefs in demon worlds. I will point out how each design account fares with respect to Cohen's argument. We will see that some avoid it and some do not.

Design Accounts of Epistemic Justification

Design accounts of epistemic justification characterize justified beliefs in terms of cognitive system design, either relative to some appropriate environment or without reference to any environment. Although this approach comes with some problematic features, it also holds out promise for yielding a determinate account of epistemic justification which is both psychologically realistic and intuitive. In this section, I will explore three alternative design accounts. Although each of these has some difficulties, they also each have enough intuitive appeal to make it worthwhile to begin to chart the "design account" territory. In general, such views provide a way of making satisfaction of the constraint of psychological realism non-accidental. This is because, assuming that the human cognitive system is designed, the view leaves the details of cognitive design to psychological investigation. I will regard the three design views which we will look at as alternative

instantiations of our schema for accounts of epistemic justification (the conjunction of (5.2) and (5.3)). I will continue to assume that this schema is both intuitive and psychologically realistic.

In addition to being potentially intuitive and psychologically realistic, design views have a good chance of being general enough to be species-neutral. That is, they are not likely to exclude from the capacity for justified belief possession any possible cognizers who are intuitively capable of possessing justified beliefs. In the case of our account schema, the psychological capacities required for justified belief are minimal. A cognizer need only possess the capacities of belief acquisition, maintenance and elimination. Design accounts add that the cognizer's cognitive system must be designed. The details of design for a particular psychological species are left unspecified. Such a specification would be needed only for the purpose of formulating the particular set of epistemic principles which applies to that psychological species. So design accounts have a kind of indexical element. The justification of a particular cognizer's beliefs is a function of the way in which the cognitive systems of cognizers of that kind are designed. This general, species-neutral feature minimizes the chances of the psychological unreality of a candidate design account, given the psychological reality of the general schema of which it is an instance.

I have suggested some general virtues of the design approach to accounts of epistemic justification. Before we proceed to a detailed scrutiny of some specific design views, we need to consider some apparent problems with any version of the design approach to epistemic justification.

Problems for the Design Approach

The first general problem which faces design accounts is that design seems to be just the sort of thing which must be eliminated from a naturalistic world view. This is because it seems to be essentially tied to purposive agency. This result may be congenial to theists, but problematic for nontheists. Can an account of epistemic justification which includes the notion of design be incorporated into a purely naturalistic world view? My response to this is two-fold. In the first place, this is a problem which many contemporary cognitive scientists have. Many cognitive scientists talk about the design of our cognitive processes without explaining how the element of design squares with a naturalistic world view. Consider the following quotation from Noam Chomsky:

The fact that all normal children acquire essentially comparable grammars of great complexity with remarkable rapidity suggests that human beings are somehow specially designed (emphasis mine) to do this, with data handling or "hypothesis-formulating" ability of unknown character and complexity (1980, p. 60).

Quotations like this can be extracted from the works of many

other leading cognitive scientists. I think that this pervasive appeal to design suggests that, in spite of the problems of squaring it with naturalism, there is something about the notion of design as an explanatory tool which it is difficult to do without.

Secondly, the problem will eventually need to be solved either by endorsing theism, giving a naturalistic account of design, or proposing intuitive, psychologically realistic accounts of epistemic justification which do not make use of the concept of design. Both the acceptance of theism and the project of providing a naturalistic account of design bring with them many problems the solutions of which are beyond the scope of this dissertation. This might make it seem most reasonable to go with the third option, that is, to look for an account of epistemic justification which does not incorporate the notion of design. However, I am not aware of any compelling argument to the effect that neither a theistic account of design nor a naturalistic account of design can be had. Furthermore, there are at least two philosophers who have recently offered accounts of epistemic justification which can be considered to be design accounts. Plantinga (1986) has proposed a design account which presupposes a theistic world view. We will look at his view and a version of his view in this section. Pollock (1986) has an account of epistemic justification which can be reasonably construed as a design account. We will look

at Pollock's view and modifications of it later as well. For these reasons, I think it is premature to dismiss design accounts out of hand.

The second general apparent problem for the design approach to epistemic justification is this. Assuming that there is either a satisfactory theistic or naturalistic account of design in the offing, is it really intuitive to say that only cognizers whose cognitive systems are designed are capable of having justified beliefs? Isn't a possible cognizer whose cognitive system has neither been created nor has been shaped by purely naturalistic evolutionary processes capable of having justified beliefs? Consider a cognizer who just comes into existence spontaneously and by pure chance who is neither created nor has evolved. Suppose that his cognitive system and the operation of his cognitive system is indistinguishable from that of a human being. Moreover, this cognitive system comes replete with appropriate memories and concepts. Suppose further that he forms the belief that there is something red before him upon being appeared to redly. Isn't it appropriate to say that this belief of his is epistemically justified? Is so, then design is not necessary for epistemic justification.

One way of defending the design approach at this point is to say that we would attribute justified beliefs to the human-like cognizer only insofar as the cognizer's beliefs are formed and maintained just as ours would be when

our cognitive systems are operating as they were designed (by creation or mere evolution) to operate. If this is right, then what is important here is not a causal connection with a designing mechanism or agent but instead a functional equivalence with some cognitive system which has been designed. However, this proposed defense misses the mark. We might be inclined to attribute justified beliefs to the human-like cognizer in circumstances in which we are ignorant of its origin. However, it seems reasonable to say that, were we to find out that the cognizer had neither evolved nor been created (by man or God), we would not know whether to attribute justified beliefs to it or not. Without knowing what something is for or how it is supposed to work, how are we to know what the functions of its cognitive states are (or for that matter, whether they are really cognitive states at all)? And if we do not know this, how are we to know whether its cognitive states are serving any end for the cognizer, epistemic or otherwise? At most, it would be reasonable to say of such a cognizer that it would have justified beliefs were its cognitive states designed to serve the roles which ours have been designed to serve. Consequently, this alleged general counterexample to the design approach to epistemic justification fails.

The final apparent general problem which we will consider before examining some specific design views is not really a potential problem for the design approach but more

one for the indispensability thesis which I defended in chapter two. Recall that this thesis holds that psychological investigation, broadly construed, is indispensable as a tool for testing accounts of epistemic justification for psychological reality. It might seem as if design accounts are a counterexample to this thesis. According to a design account, the epistemic justification of a belief is a function of the way in which a cognitive system is designed. Are not all such views automatically psychologically realistic? No. There are at least two reasons for this negative response. The first is that, as I said above, design accounts are instantiations of our account schema. Psychological investigation is required for determining whether this schema is psychologically realistic. So psychological investigation is ipso facto required for determining whether any instance of this schema is psychologically realistic. The second reason for saying that design accounts are not automatically psychologically realistic is that the thesis that human cognitive systems are designed is an empirically testable one. Psychological investigation, in the form of either introspection or scientific methodology or both, is required for confirming that the human mind exhibits design. This is not something which is a priori knowable. If it should turn out that human minds are not designed, then all design accounts are psychologically unrealistic. So design accounts do not constitute a counterexample to the indispen-

sability thesis.

It is now time to turn to a discussion of three design views. Although I endorse the general design approach, I will not defend one of these particular views against the other two. Instead, I will point out some of the strengths and weaknesses of each. Perhaps future research will lead me to a design view which is relatively unproblematic. This agnosticism on my part is in keeping with the general thrust of the dissertation and of this chapter. The overall thesis of the dissertation concerns the relation between determining the correct account of epistemic justification and psychological investigation. My thesis that one cannot succeed at the former task apart from the latter (the thesis of epistemological naturalism of chapter one) is based on the view that there is a constraint of psychological realism on accounts of epistemic justification which requires psychological investigation in its employment (the indispensability thesis of chapter two). So the constraint of psychological realism is closer to the heart of the dissertation than is the constraint of intuitiveness. This chapter is mainly intended to illustrate how research concerning the nature of epistemic justification can be guided by considerations of psychological reality. Since the three design views stand or fall together in terms of their satisfaction of the constraint of psychological realism, it is really the constraint of

intuitiveness which plays the major role in settling the dispute between them. Defending one design account as more intuitive than the other falls somewhat outside the scope of the dissertation. However, unless some design account is intuitive, the whole approach is wrong-headed. Thus, I hope that the discussion to follow gives some reason to think, not that a particular design account is intuitive, but that some design account or other is intuitive.

The first design view I will examine is a reliabilist one. Problems with this view will lead into a consideration of two "proper functioning" accounts. Each of these design account formulations will be a conjunction of (5.2), (5.3) and a proposal about what property makes belief processes the right kind for the purpose of epistemic justification.

A Reliabilist Design Account

In "Strong and Weak Justification" (unpublished), Goldman endorses a qualified form of the view that ". . . a (J-) rule system is right in (a possible world) W just in case it has a high truth-ratio in W (p. 17; the qualifications he adds to this need not concern us here). This criterion of J-rule rightness applies to Goldman's notion of strong justification. As I have noted, and as Goldman admits, this view entails that no one is strongly justified in demon worlds, as described in Cohen's counterexample. In

addition, as I have said, Goldman holds that what distinguishes the good reasoner from the bad reasoner in the demon world is that one is weakly justified and the other is not. That is, one is epistemically (merely) blameless and the other is not. Is there a version of reliabilism, other than "normal worlds" reliabilism, which does not fall victim to Cohen's counterexample? Let us consider the view I will call "design reliabilism" as a possibility.

Design reliabilism is the conjunction of (5.2), (5.3) and the following:

- (5.14) B is a belief process of the right kind if and only if B would produce a high truth-ratio of beliefs in environments of the sort for which B was designed.

This view makes the justifiedness of a belief a function of belief process reliability not in worlds of certain kinds, but in certain kinds of environments. This idea of "environments-for-which-a-process-was-designed" is familiar from animal studies. For example, certain neurons in a frog's retina seem to have been designed to trigger a feeding response on the part of the frog when the appropriate kind of stimulus is present. In the frog's normal habitat, this stimulus takes the form of an insect. However, the stimulus can be artificially simulated by experimenters using objects other than insects. Assume, what is quite unlikely, that the frog could be said to have beliefs. Then the design

reliabilist account of epistemic justification would label the frog's experimenter-induced false belief that there is an insect before it justified, since the belief would be produced by a process which would tend to produce a true belief in an appropriate environment.

Is design reliabilism weaker than strong justification? That is, does it avoid Cohen's counterexample? A cognitive process could possess the property expressed by (5.14) in environments in which the beliefs it produces are consistently false. These would be environments other than those for which the process was designed, such as the demon-world environment might be in Cohen's example. However, in spite of this, design reliabilism fails to escape Cohen's counterexample. This is because it is possible that a demon is at work in every environment of the sort for which a person's belief processes were designed. If this were the case, then none of that person's belief processes would possess the property expressed by (5.14). If none of the person's belief processes possessed this property, then, according to design reliabilism, the person would not have any epistemically justified beliefs in any worlds or environments. But this runs counter to the intuition that one who reasons well in a demon world or demon environment has justified beliefs. So, given the intuitions that drive Cohen's argument, design reliabilism fails to satisfy the intuitiveness constraint.

If we assume that the correct account of epistemic justification should avoid Cohen's counterexample, then Goldman's conception of strong justification is not adequate as an account of epistemic justification. Furthermore, although it avoids Cohen's counterexample, the normal worlds approach to reliabilism is a dubious candidate because it is fraught with difficulties. Moreover, as we have seen, design reliabilism runs afoul of the demon argument as well. Thus, given the assumption that it is possible to be epistemically justified in demon worlds, if there is hope for a version of reliabilism, it must be a version which we have not yet encountered. However, even given the denial of this assumption, there is another reason for thinking that at least no design reliabilist instantiation of our account schema ((5.2) and (5.3)) will satisfy the intuitiveness constraint (perhaps this is true of any reliabilist instantiation of this schema). This reason for thinking that it does not is that it appears to label beliefs justified which are produced by "accidentally reliable" processes. But this is counterintuitive. What follows is a purported counterexample to design reliabilism which illustrates this.

Suppose that, due to a head injury, S's visual processes produce, not only normal visual beliefs, but also beliefs about the stock market. Occasionally, when S is forming visual beliefs, S finds herself believing some proposition or other about the stock market. In addition, a

high percentage of these stock market beliefs happen to be true. Thus, S's visual processes reliably produce both normal visual beliefs and beliefs about the stock market. Furthermore, quite apart from there visually-produced stock market beliefs, S believes herself to be extremely knowledgeable about the stock market. She believes that she has developed a knack for making stock market predictions on the basis of this purported knowledge. Consequently, whenever she considers what is actually one of her visually-produced stock market beliefs, she believes that she holds it on the basis of this allegedly well-informed skill (even though she actually does not). Furthermore, she has no reasons for thinking that any of these beliefs are false. Thus, S's stock market beliefs are undefeated. Finally, S forms all of her beliefs in environments for which her belief processes were designed. Now vision is a process which has been designed if any process has. Consequently, this appears to be a case in which a process is producing a high truth-ratio of beliefs in environments for which it was designed. In fact, it is producing a high truth-ratio of beliefs of two different kinds: visual beliefs and stock market beliefs. However, whereas it is intuitive to say that the visual beliefs it produces are epistemically justified, it is not intuitive to say that the stock market beliefs it produces are epistemically justified. Hence, it is not sufficient for the epistemic justifiedness of a

belief that it be produced or sustained by a cognitive process which would be reliable in environments for which the reliable process was designed, even when the belief is defeated.

There are at least two responses which the design reliabilist could make to this alleged counterexample. The first is that the processes which produce both visual beliefs and stock market beliefs are different processes from S's pre-injury visual processes. They are processes which have been created accidentally by means of S's head injury and are consequently no longer designed processes. However, it is quite reasonable to say that S's visual processes, after her injury, are not different processes from her visual processes before the injury. Instead, they are the same processes, but they are now malfunctioning. That is, they are not functioning the way they were designed to function. The second response which the design reliabilist might make is that a different process is responsible for S's post-injury visual beliefs than is responsible for her post-injury stock market beliefs which are tied to vision. The first process is the process which takes visual percepts as input and delivers visual beliefs as output. The second process takes visual percepts as input and yields stock market beliefs as output. the problems with this response is that it is based on a misunderstanding of the process in question. It is one process which always takes

visual percepts as input, sometimes yielding only visual beliefs as output and sometimes yielding a conjunction of a visual belief with a stock market belief as output. The latter occasional result exhibits the malfunctioning of the process which is due to S's head injury.

The malfunctioning of S's visual processes is an example of a designed process which reliably produces undefeated but unjustified beliefs (the stock market beliefs) in environments for which the process was designed. The reason the reliability of the process is not sufficient to make these undefeated beliefs justified is that the reliability of the process in producing stock market beliefs is accidental reliability. What modifications can be made to (5.14) to yield an account of epistemic justification which will avoid this unintuitive result? Next we will consider a version of a proposal made by Plantinga in (1986) which purports to skirt this difficulty. This new design account makes epistemic justification a function of "proper-functioning-relative-to-an-appropriate-environment."

Proper Functioning Accounts

Plantinga's Account. Plantinga's account of epistemic justification in (1986) is as follows:

. . . a belief B has positive epistemic status for S if and only if that belief is produced in S by his epistemic faculties working properly in an appropriate environment; and B has more positive epistemic status than B* for S iff B has positive epistemic status for S and either B* does not or else S is

more strongly inclined to believe B than B* (p. 17). In considering Plantinga's view, I will ignore the biconditional after the semicolon, which is included to account for degrees of epistemic justifiedness. It is clear from the context of the article that, for Plantinga, 'positive epistemic status' is a synonym for 'epistemic justification.' Plantinga's view is on the right track as a replacement for design reliabilism. Since it analyzes epistemic justification in terms of the proper functioning of epistemic faculties, it excludes cases of accidental reliability due to the malfunctioning of a belief process. Moreover, for Plantinga, a cognitive faculty is functioning properly if and only if it is ". . . functioning in the way in which it was designed to function" (p. 14; emphasis mine). So Plantinga's position is a design account.

As it stands, Plantinga's account fails to avoid Cohen's demon world counterexample. It is easy to see that this is the case. According to Plantinga, a belief is epistemically justified only if it is produced in a person by means of the proper functioning of her cognitive faculties in an appropriate environment. By 'appropriate environment,' Plantinga appears to mean 'environment of the sort for which one's cognitive processes were designed.' But an evil demon could arrange to have a person transported to an inappropriate environment in which the person would nonetheless have entirely normal perceptual experiences. On Plan-

tinga's view, this person's perceptual beliefs would all be unjustified, even though her perceptual processes would be functioning properly. For those who share Cohen's intuitions, this is intuitively wrong. Can Plantinga's account be modified in order to avoid this result?

The following amended version of Plantinga's view appears to avoid Cohen's counterexample. This "proper functioning" view is a conjunction of (5.2), (5.3) and:

- (5.15) B is a belief process of the right kind if and only if B is functioning properly, that is, B is functioning the way it was designed to function.

Unlike Plantinga's view, this revised view does not require that the cognizer actually be in an environment of the sort for which her cognitive processes were designed in order for her to have epistemically justified beliefs. It requires merely that her processes be functioning the way they were designed to function. It might seem that this requirement could be satisfied by a person in a Cohen-style demon world. It might seem intuitive that reasoner A in Cohen's example (the good reasoner) forms, retains and eliminates beliefs as the result of the proper functioning of his belief processes. His processes are functioning as they were designed to function. On the other hand, B's processes are not functioning properly. Hence, it would seem that the modified proper functioning account sidesteps Cohen's demon world argument.

Whether or not this is the case depends on how proper functioning is analyzed. There is a reliabilist analysis of proper functioning and a non-reliabilist analysis of proper functioning. Although the "accidental reliability" counterexample to design reliabilism purports to show that reliability is not sufficient to make an undefeated belief justified, it does not show that reliability is not necessary for justified belief. So, as far as that example is concerned, reliability could be a necessary condition of proper functioning. This type of proper functioning account would be a conjunction of (5.2), (5.3), (5.15) and:

- (5.16) A belief process B is functioning properly only if B is reliable.

A reliabilist proper functioning account would not label the beliefs of demon world dwellers justified. Consequently, assuming that this is unintuitive, we must find a non-reliabilist analysis of proper functioning. Plantinga appears to suppose that proper functioning involves reliability (1986, p. 7). In an example which purports to show that "epistemic dutifulness" (a responsibility account) is not sufficient for epistemic justification, Plantinga asks the reader to imagine a person who, when appeared to redly, has a strong inclination or impulse to believe that her is not appeared to redly. Plantinga suggests that the presence of this unreliable impulse might be due to the agency of a Cartesian evil demon. His diagnosis of this situation is

that the person's cognitive faculties are not working properly. This suggests that Plantinga assumes a reliabilist account of proper functioning.

If the correct account of epistemic justification must label some demon-world beliefs epistemically justified, then we must alter Plantinga's view once again. The modification involves omitting (5.16) from the proper functioning account. Although it is not couched in terms of proper functioning, I believe that a recent view of John Pollock concerning epistemic justification in (1986) can be construed as such a non-reliabilist proper functioning position.

Pollock's Account. Pollock summarizes his general account of epistemic justification as follows: "A person's belief is justified if and only if he holds it in conformance to his epistemic norms" (p. 168). According to Pollock, ". . . our epistemic norms are the norms which actually govern our reasoning" (p. 168). This governance relation is a psychological process (p. 168). Presumably, governance by epistemic norms can take the form of either production of a belief by means of the operation of a belief acquisition process or maintenance of a belief by means of a belief maintenance process. Epistemic norms also play a corrective role. Thus, governance of reasoning by epistemic norms involves belief elimination processes. In addition, ". . . epistemic norms are to be understood in terms of pro-

cedural knowledge involving internalized rules for reasoning" (p. 132). Although possession of this procedural knowledge is sufficient for knowing how to reason, this reasoning competence need not always be exercised in actual reasoning (p. 132). This sketch of Pollock's view makes it reasonable to think that it is an instance of our account schema (the conjunction of (5.2) and (5.3)).

What then, on Pollock's view, makes some belief processes the right kind and others the wrong kind for the purpose of epistemic justification? Pollock's comment on this question is as follows:

No doubt some philosophers will be disturbed by the fact that my analysis of epistemic justification does not characterize justified beliefs in terms of a single general property (like reliability) intrinsic to the beliefs, but instead characterizes justified beliefs in terms of the reasoning underlying them. That, however, is just the way things are. What makes a belief justified is its being supported by reasoning of an approved sort, and there is no reason to think there are general intrinsic properties of beliefs that determine whether that is possible (p. 169).

Pollock denies that what distinguishes epistemically justified from epistemically unjustified beliefs is the reliability of the processes which produced them. He says instead that ". . . what makes a belief justified is its being supported by reasoning of an approved sort." This appears to be equivalent to our account schema. But this is unsatisfying. What is it that distinguishes approved reasoning from unapproved reasoning? Pollock has his own answer to this

question (see pp. 142-148). Instead of considering it, I would like to propose that Pollock's account be construed as a proper functioning view. The link is made if we say that a belief is held on the basis of the proper functioning of a person's belief processes if and only if it is held in conformance to his epistemic norms. On a purely naturalistic picture, proper functioning is a matter of evolutionary design. Thus, given this construal of Pollock's account, regulation of belief by means of epistemic norms is also a matter of evolutionary design. In this way, Pollock's account becomes a design account.

This dispels a little of the mystery concerning what it is that makes reasoning approved on a non-reliabilist proper functioning view. However, there is a fundamental problem with any non-reliabilist proper functioning view which needs to be recognized. Both theistic and naturalistic non-reliabilist proper functioning accounts have the difficulty of finding a non-circular way of distinguishing cognitive functioning which is proper for prudential ends from cognitive functioning which is proper for epistemic ends. It seems plausible to suppose that some of our cognitive processes were designed to produce mostly false beliefs in certain circumstances, since it is clear that false beliefs can have survival value. For instance, believing that a striped object one sees in the jungle is a tiger even when it is not is a process which promises to promote sur-

vival. The problem is this. Although beliefs produced by such processes may be prudentially justified, they do not seem intuitively to be epistemically justified, even on a non-reliabilist account. Given this, how can proper functioning be linked to epistemic justification?

It is interesting to note that this is a problem which does not arise for design reliabilism nor for a reliabilist proper functioning view (either environment relative or otherwise). The reason for this is that, for these views, justification-conferring processes are those which would result in mostly true beliefs in the environments for which they were designed. Such accounts label beliefs epistemically unjustified which are produced by processes which yield mostly false beliefs in environments for which they were designed. So the presence of such processes is not a problem for these reliabilist design views.

I think that the best response available for the non-reliabilist proper functioning theorist is to say that, one thing that distinguishes proper prudential functioning from proper epistemic functioning is that proper epistemic functioning involves the corrective feature of belief elimination processes, whereas proper prudential functioning does not. The operation of this defeasibility constraint can lead from the possession of an epistemically justified belief to the possession of an epistemically justified belief. Unfortunately, this is not yet a completely satis-

factory answer. What is needed in addition is a non-circular and illuminating account of epistemic defeasibility. This would be a challenging task.

Summary, Conclusion and Prospectus

This ends my illustration of the role of the constraint of psychological realism in guiding the construction of an account of epistemic justification. I have considered various instantiations of an account schema which is arguably both intuitive and psychologically realistic. Given current psychological investigation, it is reasonable to conclude that some reliabilist accounts are psychologically realistic. Whether or not reliabilist accounts satisfy the intuitiveness constraint depends on, among other things, one's conviction about the possibility of having epistemically justified beliefs in demon worlds. I considered three design accounts of epistemic justification as alternatives to Goldman's conception of strong justification as an account of epistemic justification. Design reliabilism does not avoid Cohen's counterexample. In addition, it seems to falter because of the possibility of accidentally reliable processes. This last consideration suggests a proper functioning view. Plantinga's version of this view is environment-relative and probably reliabilist. For both of these reasons, his account runs afoul of Cohen's counterexample. Pollock's "epistemic norm" account of epistemic justifica-

tion can be construed as a non-reliabilist proper functioning account. Although it avoids Cohen's counterexample, one problem with this view, which is not shared by the reliabilist views, is that it may not enable one to distinguish functioning which is proper for prudential purposes from functioning which is proper for epistemic purposes. A defeasibility constraint may be a way out of this difficulty.

In this dissertation, I have argued for the thesis of epistemological naturalism, which is the view that it is not possible to determine the correct account of epistemic justification apart from psychological investigation, broadly-construed. The argument for this thesis is based on the constraint of psychological realism and what I have called the "indispensability thesis." According to this constraint, no account of epistemic justification is correct which entails that humans are not capable of possessing justified beliefs. The indispensability thesis is the view that an account of epistemic justification cannot be tested for psychological reality apart from psychological investigation. I also argued that the constraint of psychological realism has an important role to play both negatively, in terms of testing proposed accounts of epistemic justification for psychological reality, and positively, in terms of guiding the construction of such an account. Moreover, I provided illustrations of both of these roles of the con-

straint.

The thesis of epistemological naturalism has two sub-theses. One of these, the "replacement thesis," is the view that psychological investigation is both necessary and sufficient for determining the correct account of epistemic justification. I argued that, although psychological investigation is necessary for this purpose, it is not sufficient. This view is the "dependency thesis," which is the other sub-thesis of the thesis of epistemological naturalism. Since psychological investigation is not sufficient for determining the nature of epistemic justification, I endorsed an additional constraint, the intuitiveness constraint, satisfaction of which is a necessary condition for determining the correct account of epistemic justification. These two constraints may be both individually necessary and jointly sufficient for determining the nature of epistemic justification.

Given these conclusions, what are some directions for future research? In the first place, similar constraints of psychological realism could be employed to test and construct theories of epistemic justification which are potential answers to questions (1.1) and (1.2). Recall that question (1.1) concerns the structure of epistemic justification and question (1.2) concerns the specific conditions under which human beliefs of various kinds are epistemically justified. Answers to this latter question would take the

form of lists of epistemic principles. Secondly, versions of the constraint of psychological realism could be used to test, not only accounts of epistemic justification and theories of rationality, but also accounts of propositional knowledge, given a constraint generated by the denial of skepticism with respect to knowledge. It is interesting to note in this connection that the constraint of psychological realism, as defined in chapter two, applies to accounts of knowledge which include epistemic justification as a necessary condition. Since, on these views, knowledge requires justification, we cannot determine fully what knowledge is without determining what epistemic justification is. Since we cannot determine the correct account of epistemic justification without determining whether it is psychologically realistic, it follows that it is not possible to know what knowledge is without knowing of a given account of knowledge whether it is psychologically realistic. Furthermore, given that we cannot know what epistemic justification is apart from psychological investigation (this is the thesis of epistemological naturalism), it can be concluded that we cannot know what knowledge is apart from psychological investigation. All of the considerations adduced in this paragraph reveal that the conclusions of this dissertation have ramifications much wider than the task of determining the nature of epistemic justification.

Another future project which needs to be pursued,

given the conclusions of this dissertation, is to attempt to spell out the alternative implications of the thesis of epistemological naturalism which were mentioned at the end of chapter two. Recall that, if this thesis is true, then either giving an account of epistemic justification is not just to analyze the concept of epistemic justification, or conceptual analysis is an activity which is not entirely a priori. If either one of these positions is correct, then there will need to be a considerable shift away from the traditional understanding of the nature of epistemology. Related to this question is the question concerning whether satisfaction of the constraint of intuitiveness and the constraint of psychological realism are not only individually necessary, but also jointly sufficient, for determining the correct account of epistemic justification.

Finally, the whole purpose of our discussion is to contribute, eventually, to the determination of the correct account of epistemic justification. The full development of such an account in light of our conclusions is another task which remains undone. I have proposed that design accounts be considered as potential candidates. Although these types of accounts seem worth pursuing as intuitive and psychologically realistic views, if they ultimately fail, we must seek another type of position which promises to satisfy the two constraints we have endorsed. Future research should be aimed at finding such an account.

APPENDIX: HARMAN ON BAYESIAN EPISTEMOLOGY

In (1986, pp. 25-27), Gilbert Harman rejects the following epistemological claim:

- (A.1) The updating of (subjective) probabilities via conditionalization or generalized conditionalization is (or ought to be) the only principle of reasoned revision (of degrees of belief).

Harman's reasons for dismissing this claim are reminiscent of arguments which employ the constraint of psychological realism. An abbreviated version of his argument against (A.1) is as follows: "The probability assignments required for conditionalization lead to a combinatorial explosion. Consequently, it is not possible to represent these assignments explicitly. The only way to represent them implicitly is by means of a general principle. However, on this approach, neither conditionalization nor generalized conditionalization is usable by humans. This is because conditionalization requires that one remember all the immediate perceptual evidence that has influenced one's present degrees of belief. But humans rarely remember such evidence beyond the moment in which it is originally acquired. Furthermore, although generalized conditionalization does not have this requirement, it leads to a combinatorial explosion."

Harman's strategy here is to point out the psychological components of (A.1) and to argue for the claim that humans either do not have the requisite psychological capacities or do not exercise these capacities. Does he adduce any psychological evidence to support this claim? In the context of a similar argument in (1985, p. 244) he appeals to Kahneman, Slovic and Tversky (1982) and Nisbett and Ross (1980) for evidence that humans do not reason well probabilistically. So Harman does employ the results of some scientific psychological investigation to make his case. However, in the above argument, explicit appeal to psychological evidence is lacking. The two psychological claims on which Harman bases his claim that the principle referred to by (A.1) is unusable by humans are: (a) humans rarely remember immediate perceptual evidence beyond the moment in which it is acquired and (b) humans are not capable of keeping track of a combinatorial explosion of representations of probability assignments. It appears that Harman implicitly assumes that these claims can be confirmed on the basis of introspection.

Harman concludes, given these considerations, that the principle to which (A.1) refers is not a principle which is actually used by humans. Does he conclude that this principle is not normatively correct (i.e., does he conclude that it is not the case that humans ought to reason in accord with this principle)? This is not clear. Of

course, given Harman's argument and the eminently plausible thesis that humans actually possess reasonable beliefs, it can be concluded that it is not the case that humans ought to reason by means of conditionalization or generalized conditionalization. This added explicit rejection of skepticism makes Harman's argument an application of a constraint of psychological realism to (A.1).

REFERENCES

- Alston, William P. (1976). "Two Types of Foundationalism," Journal of Philosophy 73:165-185.
- Alston, William P. (1985). "Concepts of Epistemic Justification," The Monist 68:57-89.
- Alston, William P. (Unpublished). "Justification and Knowledge."
- Anderson, John R. (1984). "Cognitive Psychology," Artificial Intelligence 23:1-11.
- Audi, Robert (1980). "Foundationalism and Epistemic Dependence," Journal of Philosophy 77:612-613.
- Block, Ned, ed., (1980) Readings in Philosophy of Psychology, vol. 1 (Cambridge, Mass.: Harvard University Press).
- Bonjour, Laurence (1980). "Externalist Theories of Empirical Knowledge," Midwest Studies in Philosophy 5:53-73.
- Bonjour, Laurence (1985). The Structure of Empirical Knowledge (Cambridge, Mass.: Harvard University Press).
- Chaplin, J. P. (1975). Dictionary of Psychology, second edition (New York, New York: Dell Publishing Co.).
- Cherniak, Christopher (1983). "Rationality and the Structure of Human Memory," Synthese 57:163-186.
- Cherniak, Christopher (1986). Minimal Rationality (Cambridge, Mass.: The MIT Press).
- Chisholm, Roderick M. (1977). Theory of Knowledge, second edition (Englewood Cliffs, New Jersey: Prentice-Hall, Inc.).
- Chomsky, Noam (1959). "A Review of B. F. Skinner's Verbal Behavior" Language 35:26-58. Reprinted in condensed form in Block (1980), pp. 48-63.
- Clarke, Murray (1986). "Doxastic Voluntarism and Forced Belief," Philosophical Studies 50:39-51.

- Cohen, Stewart (1984). "Justification and Truth," Philosophical Studies 46:279-295.
- Dretske, Fred I. (1981). Knowledge and the Flow of Information (Cambridge, Mass.: The MIT Press).
- Firth, Roderick (1978). "Are Epistemic Concepts Reducible to Ethical Concepts?," in Values and Morals, Alvin I. Goldman and Jaegwon Kim, eds. (Dordrecht, Holland: Reidel).
- Fodor, Jerry (1983). The Modularity of Mind (Cambridge, Mass.: The MIT Press).
- Gettier, Edmund (1963). "Is Justified True Belief Knowledge?," Analysis 23:121-123.
- Goldman, Alvin I. (1978). "Epistemics: The Regulative Theory of Cognition," Journal of Philosophy 75:509-523.
- Goldman, Alvin I. (1979). "What is Justified Belief?," in Justification and Knowledge, George S. Pappas, ed. (Dordrecht, Holland: Reidel)
- Goldman, Alvin I. (1985). "The Relation Between Epistemology and Psychology," Synthese 64:29-68.
- Goldman, Alvin I. (1986). Epistemology and Cognition (Cambridge: Harvard University Press).
- Goldman, Alvin I. (Unpublished). "Strong and Weak Justification."
- Goldman, Alvin I. and Kim, Jaegwon, eds. (1978). Values and Morals (Dordrecht, Holland: Reidel).
- Harman, Gilbert (1973). Thought (Princeton, New Jersey: Princeton University Press).
- Harman, Gilbert (1985). "Positive versus Negative Undermining in Belief Revision," in Kornblith (1985a), pp. 231-248.
- Harman, Gilbert (1986). Change in View (Cambridge, Mass.: The MIT Press).
- Heil, John (1983a). "Believing What One Ought," Journal of Philosophy 80:752-765.
- Heil, John (1983b). "Doxastic Agency," Philosophical Studies 43:355-364.

- Kahneman, Daniel, Slovic, Paul and Tversky, Amos, eds. (1982). Judgment Under Uncertainty: Heuristics and Biases (Cambridge: Cambridge University Press).
- Kornblith, Hilary (1982). "The Psychological Turn," Australasian Journal of Philosophy 60:238-253.
- Kornblith, Hilary (1983). "Justified Belief and Epistemically Responsible Action," The Philosophical Review 92:33-48.
- Kornblith, Hilary, ed. (1985a). Naturalizing Epistemology (Cambridge, Mass.: The MIT Press).
- Kornblith, Hilary (1985b). "Introduction: What is Naturalistic Epistemology?," in Kornblith (1985a), 1-13.
- Kornblith, Hilary (Unpublished). "Introspection and Misdirection."
- Lehrer, Keith (1974). Knowledge (Oxford: Clarendon Press).
- Lehrer, Keith (1986). "The Coherence Theory of Knowledge," Philosophical Topics 14:5-25.
- Lehrer, Keith and Cohen, Stewart (1983). "Justification, Truth and Coherence," Synthese 55:191-207.
- Miller, George (1956). "The Magical Number Seven Plus or Minus Two: Some Limits of Our Capacity for Processing Information," Psychological Review 63:81-97.
- Moore, G. E. (1922). Philosophical Studies (London: Routledge and Kegan Paul, Ltd.).
- Moore, G. E. (1959). "Proof of an External World," in Philosophical Papers (London: Allen and Unwin).
- Moser, Paul K. (1985). Empirical Justification (Dordrecht: D. Reidel Publishing Co.)
- Nisbett, Richard and Ross, Lee (1980). Human Inference: Strategies and Shortcomings of Social Judgment (Englewood Cliffs, New Jersey: Prentice-Hall, Inc.)
- Nozick, Robert (1981). Philosophical Explanations (Cambridge, Mass.: The Belknap Press of Harvard University Press).
- Pappas, George S., ed. (1979a). Justification and Knowledge

(Dordrecht: D. Reidel Publishing Co.)

- Pappas, George S. (1979b). "Basing Relations," in Pappas (1979a).
- Plantinga, Alvin (1986). "Epistemic Justification," Nous 20:3-18.
- Pollock, John (1974). Knowledge and Justification (Princeton: Princeton University Press).
- Pollock, John L. (1979). "A Plethora of Epistemological Theories," in Pappas (1979a), 93-113.
- Pollock, John L. (1986). Contemporary Theories of Knowledge (Totowa, New Jersey: Rowman and Littlefield).
- Price, H. H. (1954). "Belief and Will," Proceedings of the Aristotelian Society Supplementary Volume 54:1-22.
- Quine, Willard van Orman (1960). Word and Object (Cambridge, Mass.: Technology Press of MIT).
- Quine, Willard van Orman (1961). "Two Dogmas of Empiricism," in From a Logical Point of View (Cambridge, Mass.: Harvard University Press).
- Quine, Willard van Orman (1969). "Epistemology Naturalized," in Ontological Relativity and Other Essays (New York: Columbia University Press), 69-90 (Reprinted in Kornblith 1985a, pp. 15-29).
- Reid, Thomas (1970). An Inquiry into the Human Mind (Chicago: University of Chicago Press).
- Senor, Thomas D. (Unpublished). "Justification, Truth, and Epistemological Assumptions."
- Sosa, Ernest (1980). "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge," in Midwest Studies in Philosophy 5:3-25.
- Swain, Marshall (1979). "Justification and the Basis of Belief," in Pappas (1979a), 25-49.
- Swinburne, Richard (1981). Faith and Reason (Oxford: Oxford University Press).
- Tulving, E. (1983). Elements of Episodic Memory (Oxford: Oxford University Press).

- Tversky, Amos and Daniel Kahneman (1982). "Judgment Under Uncertainty: Heuristics and Biases," in Kahneman, Slovic and Tversky, eds. (1982).
- Wason, Peter C. and Philip N. Johnson-Laird (1972). The Psychology of Reasoning: Structure and Content (London: Batsford).
- Wickelgren, W. A. (1976). "Memory Storage Dynamics," in W.K. Estes, ed. Handbook of Learning and Cognitive Processes 4 (Hillsdale, New Jersey: Erlbaum).
- Williams, Bernard (1973a). Problems of the Self (Cambridge: Cambridge University Press).
- Williams, Bernard (1973b). "Deciding to Believe," in Williams (1973a).
- Winters, Barbara (1979) "Believing at Will," Journal of Philosophy 76:243-256.
- Wolterstorff, Nicholas (Unpublished). "Belief, the Will and Noetic Obligations."