

## **INFORMATION TO USERS**

**This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.**

**The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.**

**In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.**

**Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.**

**Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.**

# **UMI**

**A Bell & Howell Information Company  
300 North Zeeb Road, Ann Arbor MI 48106-1346 USA  
313/761-4700 800/521-0600**



**Inferential Justification**

by

**Stephen Andrew Fogdall**

A dissertation submitted in partial fulfillment  
of the requirements for the degree of

Doctor of Philosophy

University of Washington

1997

Approved by William J. Talbot  
Chairperson of Supervisory Committee

Program Authorized  
to Offer Degree Department of Philosophy

Date May 24, 1997

**UMI Number: 9736271**

**Copyright 1997 by  
Fogdall, Stephen Andrew**

**All rights reserved.**

---

**UMI Microform 9736271  
Copyright 1997, by UMI Company. All rights reserved.**

**This microform edition is protected against unauthorized  
copying under Title 17, United States Code.**

---

**UMI**  
300 North Zeeb Road  
Ann Arbor, MI 48103

© Copyright 1997

Stephen Andrew Fogdall

In presenting this dissertation in partial fulfillment of the requirements for the Doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to University Microfilms, 1490 Eisenhower Place, P.O. Box 975, Ann Arbor, MI 48106, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature *Stephen G. Johnson*

Date May 24<sup>th</sup> 1997

University of Washington

Abstract

Inferential Justification

by Stephen Andrew Fogdall

Chairperson of the Supervisory Committee  
Associate Professor William J. Talbott  
Department of Philosophy

I aim to construct a partial account of *inferential epistemic justification*, the process by which epistemic justification is “transferred” from an evidentiary belief (or a set of such beliefs) to some further belief via inferential connections. The account focuses on two necessary conditions, similar versions of which have been widely discussed in the epistemological literature. The bulk of my dissertation is devoted to providing more detailed and exact formulations of these conditions than have previously been offered. Each is presented in an intuitive and rough formulation, meant to capture the spirit of the condition as discussed in the literature, and is then honed into a more precise and acceptable shape. The first condition, which I call the *Basis Condition*, requires in its rough and initial formulation that a subject’s belief be correctly inferentially based on its evidentiary beliefs if the latter are to confer justification on it; the condition requires an actual, psychologically instantiated inference process which generates the belief from its evidence. I thus agree with recent “naturalized epistemologists” who argue that mere *possession* of evidence which supports a belief is not sufficient for that belief to be inferentially justified; one must actually use the evidence in forming the belief in question. I call the second condition which I take to be necessary for inferential justification the *Access Condition*. It requires of a subject that she not simply hold her belief on the basis of adequate evidence, but that she be able to tell on reflection *that* her belief is *evidentially supported* and *causally sustained* by that evidence. Antecedents of

the Access Condition have been championed by so-called *internalist* epistemologists who hold that a belief is epistemically justified only if at least some of the factors responsible for its justification are available or accessible to the subject on reflection. The account of inferential justification which I offer can therefore be thought of as an attempt to make sense of internalism within the framework of naturalized epistemology.

## Acknowledgments

First and foremost, I wish to thank my parents and siblings. Secondly, I would like to express my gratitude to three individuals who have helped me immeasurably over the past seven years and who serve for me as models both philosophically and personally. They are Bill Talbott, Larry Bonjour, and Cass Weller. I cannot overstate the impact each of them has had on my approach to philosophical problems, and to the matters dealt with in the present work.

I would also like to acknowledge Ann Baker, Audre Brokes, Andrew Jeffrey, and Hilary Kornblith for helping me to think about issues connected with inferential justification. Finally, I must thank my dear friend Judith Weinstock. I could not have finished this dissertation without her support and enlightening advice.

## Table of Contents

One	Introduction	1
	1. <i>Inferential Justification</i>	1
	2. <i>A First Step</i>	5
	3. <i>The Secure Position</i>	14
	4. <i>Summary and Overview</i>	27
Two	Outline of a Theory	31
	1. <i>The Basis Condition</i>	32
	2. <i>The Access Condition</i>	43
	3. <i>The Question of Further Conditions</i>	56
	4. <i>Summary</i>	61
Three	The Basis Condition	62
	1. <i>Some Preliminaries</i>	62
	2. <i>Is Reasoning Governed By Rules?</i>	68
	3. <i>The Disposition-Based Account of Inference</i>	76
	4. <i>The Role of Connecting Beliefs</i>	87
	5. <i>Blind and Mediated Tracking</i>	93
	6. <i>The Basis Condition Restated</i>	104
	7. <i>Summary</i>	106
Four	The Access Condition	107
	1. <i>Overview</i>	107
	2. <i>Reflective Access and Noninferential Knowledge</i>	109
	3. <i>First and Second Order Internalism</i>	113
	4. <i>Refining the Access Condition</i>	126
	5. <i>Summary</i>	132
Five	Responses to Objections	134
	1. <i>The Theory Reviewed</i>	134
	2. <i>Is The Account Satisfiable In Principle?</i>	140
	3. <i>Is The Account Satisfied In Practice?</i>	162
	4. <i>Summary</i>	175
	Bibliography	176

# One

## Introduction

### *1. Inferential Justification.*

It goes without saying that when one comes to have a new belief, there is frequently some importance placed on having evidence for that belief. Other people may request such evidence, and one must be in a position to give it. Very often, this evidence takes the form of *other* beliefs from which one *infers* the belief in question. When a belief is supported in this way by evidence, that is, when it has been correctly inferred from an appropriate set of evidentiary beliefs, it is in a certain sense *warranted* or *justified*. This inferential justification, as it may be called, is the focus of the present work. In the chapters that follow, I will develop a partial theory of inferential justification by presenting and refining formulations of two of its necessary conditions.<sup>1</sup> The rationale for the theory, partial though it is, will prove to be that it illuminates and clarifies something which, at least much of the time, we value quite highly, namely, having *reasons* or *evidence* for our beliefs which put us in a *secure position* with respect to the truth.

It is important to be clear on the sort of justification that is at issue here. There are many senses in which beliefs can be said to be justified or unjustified. For example, a belief might be said to be justified because it is *prudent*, as in the case, we may pretend, of an athlete who, struggling against great odds to break a long standing track and field record, must believe that she will break the record in order to have any significant chance of doing so. What warrants or justifies her belief in this circumstance is the fact that such belief is necessary for achieving her goal. But her belief may not be justified in the sense

---

<sup>1</sup>I will also try to indicate, where I can, what conditions could be added to complete the account.

in which I am interested, for she may not, and perhaps most likely *does* not, possess appropriate evidentiary beliefs from which she has inferred, or could infer, that she will break the record.

Inferential justification, the item which is my target, may be distinguished from other sorts of justification, such as the “prudential” variety just discussed, by saying that the former is a kind of *epistemic* justification, i.e., it is one species of the sort of justification which is involved in having *knowledge*. Thus, in the case just considered, the athlete could be said to *know* that she will break the record only if she were inferentially justified in believing that she will break the record. But the prudential justification that in fact attaches to her belief is not relevant to the question whether she knows this. Let us explore this point for a moment.

The question before us concerns the conditions which would have to be met in order to know *that* something is the case. Let us ask, then, what requirements I would have to fulfill in order to know that, say, there is a copy of Arthur Koestler's *The Sleepwalkers* on the desk in the next room. Knowing that *The Sleepwalkers* is on the desk in the next room requires, first of all, that it be *true* that *The Sleepwalkers* is on the desk in the next room. In other words, the *proposition* that *The Sleepwalkers* is on the desk in the next room must be true, in the sense of corresponding with objective, mind-independent reality.<sup>2</sup> The second thing that knowing requires is that this true proposition, or the state of affairs that makes it true, be somehow reflected or represented *in my mind*.

---

<sup>2</sup> I shall speak freely of propositions in this dissertation without worrying about their ontological status. I shall also not worry about the fact that most of the semantic items which I will identify as propositions are not, in actuality, propositions in the strict classical sense, that is, items which are free of indexicals and timelessly true. Finally, I will assume without question the realist's correspondence theory of truth. The nature of truth is not at issue in the discussions that follow.

I must, that is, believe that *The Sleepwalkers* is on the desk in the next room; I must accept that proposition.<sup>3</sup>

Of course, the existence of a mere true belief is not by itself sufficient for having knowledge. Imagine that I am a gambler who has become so deranged that I have absolute faith in a quasi-mystical numerological technique of my own devising. As a result of applying this technique, I come to believe that the horse Sterling Shadow will win the race tomorrow, and bet accordingly. Imagine, moreover, that this belief is true, because someone (unbeknownst to me) has been hired to fix the race. It seems, nevertheless, that I do not *know* that Sterling Shadow will win.

To fill out the case a little, suppose that my numerological technique is astonishingly unreliable. Most of the beliefs which it has led me to adopt in the past, or which it will lead me to adopt in the future, have been, or will be, false. Suppose in addition, that I am well aware of my technique's unreliability, in that I remember clearly many of the instances in which it lead me astray, and would admit, if asked, that it frequently yields inaccurate results. It is thus, in some sense, unreasonable or irrational for me to believe that Sterling Shadow will win the race, given the manner in which I arrived at that belief. One way to put this is to say that my belief is *not justified* by my evidence, and in view of that fact, seems not to constitute knowledge, despite being true.

So, the question whether I know that *The Sleepwalkers* is on the desk in the next room, assuming that this proposition is true, and that I believe it, turns on whether this belief is justified in the right way. I must not be "overstepping my bounds," or "jumping to a conclusion" in holding this belief. I must have adequate evidence, or grounds, or reasons for holding it. This *epistemic* variety of justification, so called to indicate that it is the sort of justification required for knowledge, is precisely what is

---

<sup>3</sup>It is sometimes said, in recognition that beliefs can be held more or less strongly, that I must *confidently* believe a proposition in order to know that it is true. For simplicity, I will omit this modifier.

missing in the case of my belief that Sterling Shadow will win. I really ought not to hold this other belief, since I have no reason to think that it is true. Similarly, while the track and field athlete ought *in one sense* to believe that she will break the record, there is another important sense in which she ought not to believe this, since it is not justified by her evidence.<sup>4</sup>

If it is evidence or justification that I need in order to know that *The Sleepwalkers* is on the desk in the next room, where could this evidence come from? The book is not currently in my line of sight, or in other ways available to my sensory faculties. If the book were on the desk in front of me, then I might have the so-called “evidence of the senses” that it sits before me. For it is plausible that in such circumstances, when one has the visual experience or impression of the sort that is typically produced when one’s visual faculties are working properly, one is in good light, and a book is in one’s line of sight at close range, and one bases on this experience or impression a belief that there is a book before one, then this belief is *directly* or *immediately* justified by the visual experience or impression in question.<sup>5</sup>

Such direct or immediate justification, which derives from basing a belief on the appropriate sensory impression or experience, could not be involved in my knowing that

---

<sup>4</sup>The view of knowledge as epistemically justified true belief is sometimes referred to as the traditional conception (or analysis, or definition) of knowledge. It is now generally acknowledged that in light of certain counter-examples due to Gettier (1963) and subsequent theorists, the traditional conception is flawed because justified true belief is not sufficient for knowledge. Some, such as Goldman (1967) have gone so far as to suggest that the traditional conception ought to be scrapped altogether, by replacing the justification condition with a condition to the effect that the belief be produced in the right way (e.g., by being caused by the state of affairs which makes it true). Alvin Plantinga (1993a and b) has argued that epistemologists focus their attention not on justification but on *warrant*, the term he uses for that factor or quantity which *is* sufficient to convert true belief into knowledge, and which therefore avoids Gettier’s counter-examples to the traditional conception. For my part, I am still interested in justification. So long as it is allowed that, in at least some instances, something like inferential justification is *necessary* for knowledge, the investigations of this dissertation will have some point.

<sup>5</sup>Audi (1988) and (1993) defends a view of this sort, as does Pollock (1986), and Alston (1989) and (1992).

*The Sleepwalkers* is on the desk in the next room. In this case, justification derives instead from the fact that my belief about the *The Sleepwalkers* was *inferred* from certain other of my beliefs, some of which, of course, may themselves be *directly* justified. My belief that *The Sleepwalkers* is on the desk in the next room is *inferentially* justified because I correctly reasoned from beliefs of the following sort: (i) that the book was on the desk in the next room a short time ago, and (ii) that nothing has occurred in the meantime which would have caused it now to be in a different location. (There may, of course, be more general beliefs operative in my reasoning, principles pertaining to inertia and object constancy, such as, for example, that books, and other items, do not tend to randomly change location or wink out of existence. Such possible additional principles may safely be ignored for the moment.)<sup>6</sup>

The *evidence* for my belief that *The Sleepwalkers* is on the desk in the next room is thus not the evidence of the senses, but rather evidence supplied by other beliefs that I hold. Speaking intuitively, we might say that the former belief is justified *by way of* the latter evidentiary beliefs. If these evidentiary beliefs are justified, then that justification may be “transferred” by way of inference to my belief that *The Sleepwalkers* is on the desk in the next room. The project of this dissertation is to spell out some of the factors involved in this so-called transfer of justification.

## 2. A First Step.

Our starting point for the investigations in this dissertation is the observation that a belief is inferentially justified by certain other beliefs only if the latter provide genuinely good evidence for the former. Our task will be to determine some of the further

---

<sup>6</sup>One could try to suggest that my belief that *The Sleepwalkers* is on the desk in the next room is directly justified by a “memory impression” of having left it there a short time ago (see, e.g. Audi 1993, p. 309). But the most that such a memory impression could justify is the first of the premises mentioned just now. This by itself is not sufficient to justify my belief that *The Sleepwalkers* is *still* on the desk in the next room.

conditions which must be satisfied for inferential justification to be possible. I will begin to sketch these further conditions in the next chapter. First, however, we ought to focus briefly on the question of what is involved in one belief, or set beliefs, providing genuinely good evidence for another.

If we examine the case discussed in the last section, in which I inferred that *The Sleepwalkers* is on the desk in the next room from the evidence consisting of (i) my belief that *The Sleepwalkers* was in that location a short time ago, and (ii) my belief that nothing has occurred in the meantime which would cause its location to have changed, it seems that the latter have a kind of appropriateness or relevance to the former which warrants or justifies my inference. We might say, following Earl Conee and Richard Feldman, that my belief “fits” its evidence, and that it is in virtue of this “fittingness” that the latter provide good evidence for the former.<sup>7</sup>

The suggestion is that there is a potential *relation* among beliefs which can be invoked to explain why one belief, or set of beliefs, provide good evidence for another. Indeed, many epistemologists do speak explicitly of a relation of *evidential support* in this connection.<sup>8</sup> Other epistemologists use similar terms. Some speak of an “epistemic linkage” between evidence and belief,<sup>9</sup> or an “evidential tie.”<sup>10</sup> While it may be very difficult, on a philosophical level, to say what this relation involves, and why it holds or does not hold in a given case, we appear to have a broad range of intuitions concerning it. There is a large class of cases in which we have very robust and unequivocal intuitions that a certain belief, or set of beliefs, evidentially supports some further belief. Consider the following handful of somewhat banal examples.

---

<sup>7</sup>Conee and Feldman (1985).

<sup>8</sup>See, for example, Audi (1993), Feldman (1992), Moser (1989), and Plantinga (1993b).

<sup>9</sup>For example, Conee (1988).

<sup>10</sup>Nozick (1981).

First. Suppose it is the night before an election, and I believe that all of the major opinion polls show candidate X ahead of candidate Y by amounts far greater than the margins of error of the polls. Intuitively, my belief that candidate X will win the election fits my evidentiary belief that she leads in all of the major polls. The latter evidentially supports the former.

Second. Suppose that I believe that a large number of crows have been observed by myself and others, in a wide variety of circumstances and geographical locations, and that all such observed crows are black. Intuitively, my further belief that all crows are black is evidentially supported by this belief.

Third. Suppose that I hear a crash in the kitchen, and upon investigating, see the cat lapping up milk from a spilled carton on the floor. Here, my evidentiary beliefs (which most likely include (i) the belief that the carton of milk has been spilled, (ii) the belief that this could be explained by the cat's having knocked the carton over, and (iii) the belief that no other explanations appear to be available) support the belief that the cat knocked the carton over.

Fourth. Suppose that I believe that all humans are mortal, and that Timothy Leary is human. Intuitively, these beliefs evidentially support my further belief that Timothy Leary is mortal.

If we let the foregoing cases (along with similar examples that might be enumerated) provide content for talk of a relation of "evidential support" among beliefs, then we can formulate our first observation concerning inferential justification in the following way:

(IJ1) S's belief that p is inferentially justified by her beliefs  $B_1$  through  $B_n$  only if her beliefs  $B_1$  through  $B_n$  evidentially support her belief that p.

This requires, however, that we face the following worry. Focus, for a moment, on the fourth case discussed just now, involving my belief that Timothy Leary is mortal. Here it seems that the reason why my evidentiary beliefs support this belief is that they *entail* it,

in the sense that it is logically impossible for all of the former to be true if the latter is false. But, in the first three cases, the evidentiary beliefs fall short of entailing the belief they support. This suggests that relations of evidential support “come in degrees,” with entailment representing the limiting case, a suggestion which is borne out in our intuitions about particular cases. For instance, it intuitively seems that the degree of evidential support involved in the first example above is less than that involved in the third example, which is in turn less than that involved in the second. Now, if evidential support “comes in degrees,” it seems reasonable to conclude that the notion of inferential justification is itself inherently gradational, for it appears that a belief can be more or less justified depending on the quality and amount of evidence supporting it. One might therefore consider it a somewhat artificial feature of (IJ1) that it focuses on a *categorical* or *absolute* notion of inferential justification. It specifies a necessary condition of a belief’s being inferentially justified *simpliciter*, not justified to this or that degree, and it relates this notion to a similarly categorical or absolute notion of evidential support.

But any worry about the artificiality of (IJ1) would be misplaced, I believe. For there is a useful categorical or absolute notion of inferential justification, namely, the notion of inferential justification *to the degree* required for inferential *knowledge*. Similarly, there is a useful categorical notion of evidential support, namely, the notion of evidential support *to the degree* required for a belief to be inferentially justified by its evidence (i.e., the degree of evidential support required for a belief to qualify as inferential knowledge).<sup>11</sup> Of course, it would be difficult, if not impossible, to precisely specify the degree of evidential support involved in the categorical or absolute notion of inferential justification which is in question in (IJ1), and I shall make no attempt to offer a specification. But this should not be thought to indicate that the categorical notion of

---

<sup>11</sup>In speaking of inferential justification, or evidential support, *to the degree* required for knowledge, I do not mean to imply that these notions include conditions which rule out Gettier-style counter-examples to the traditional analysis of knowledge.

inferential justification is less clearly defined than the gradational notion. For we shall have just as difficult a time sorting out degrees of inferential justification or evidential support in many cases as we have in determining whether a belief is inferentially justified or evidentially supported simpliciter. What, for example, is the degree to which my evidence supports my belief that Bill Clinton is the current president of the United States, and is that degree greater than, lesser than, or equal to, the degree to which my belief that all crows are black is supported by its evidence? Our intuitions, I believe, are completely silent on this question. We will thus be no worse off if we continue to work with categorical notions of inferential justification and evidential support, as in (IJ1).

In this dissertation, I will not try to answer the question of what evidential support consists in, of what *makes* one belief, or set of beliefs, evidentially support another, for the part of the theory of inferential justification that I will defend in the following chapters can be developed by taking the notion of evidential support as a primitive. I will therefore limit myself to the following two claims concerning its ultimate philosophical analysis.

First, evidential support is *objectively truth-linked*, i.e., there is some intimate relationship between a belief's being evidentially supported and its being very probably true.<sup>12</sup> This can be seen from the following considerations: Intuitively, the difference between an item of knowledge and a merely lucky guess, as when I happen to arrive at a true belief concerning the outcome of a horse race by applying some arbitrary and

---

<sup>12</sup>Whether the degree of support which a belief receives from its evidentiary beliefs can simply be *identified* with the objective probability of the former's being true, given that the latter are true, is far from clear. For instance, on some formulations of the probability calculus, it is an axiom that if a proposition  $p$  entails a proposition  $q$ , then the probability of  $q$  given  $p$  is 1, i.e., it is maximal (see, e.g., Burks 1977). If this axiom is true, then the probability of Bill Clinton's being President, given that  $2+2=5$ , is 1. But it is not plausible to think that the proposition that  $2+2=5$  provides maximal evidential support for the proposition that Bill Clinton is President. Our purposes do not, however, require that we determine the precise relationship between evidential support and objective probability. It is enough to recognize that there is *some* intimate relationship between these quantities. (See Plantinga 1993b, chs. 8 and 9, for an extended discussion of the problematic relationship between evidential support and objective probability.)

unreliable method, is that the latter is true only by the happiest of accidents. If, on the other hand, I *know* that the horse in question is going to win (because, e.g., I know that the race is fixed), then my belief is in some intuitive sense *not* true by accident (though in what way exactly is difficult to make precise). How does an item of knowledge obtain this status of being nonaccidentally true? One way for it to achieve this status, it seems, is by being inferentially justified.<sup>13</sup> Thus, in normal circumstances, an inferentially justified belief is, if true, then nonaccidentally true, and if false, then nonaccidentally false. The best way to make sense of this is to say that inferential justification is objectively truth-linked, that an inferentially justified belief is, in normal circumstances, very probably true, and that this connection between inferential justification and objective probability is by way of the relation of evidential support (i.e., an inferentially justified belief is very probably true *because* it is supported by adequate evidence). Accordingly, it seems that evidential support is objectively truth-linked as well.<sup>14</sup>

Secondly, any of the main accounts of objective probability, and thus any plausible account of the objectively truth-linked character of evidential support, is compatible with the theory of inferential justification that I will develop.<sup>15</sup> Let me briefly sketch two of these alternative approaches in order to indicate the range of possibilities that are open.

One approach would be to develop the analogy between evidential support and entailment. The latter is a logical relation, that is, it holds (or fails to hold) among beliefs in virtue of the contents of those beliefs, and it holds *necessarily* when it holds at all.

---

<sup>13</sup>Presumably another way is for it to be directly or immediately justified.

<sup>14</sup>Actually, we have to be careful here, since inferential justification is not *sufficient* to make a true belief nonaccidentally true in the way required for knowledge (as Gettier's counter-examples show). A belief's being inferentially justified can nevertheless be an important part of the reason why it is nonaccidentally true in the way required for knowledge, and this is enough to suggest an intimate connection between inferential justification/evidential support and objective probability.

<sup>15</sup>I will defend this claim in Chapter 5.

Accordingly, one could say that evidential support, even in cases where it falls short of entailment, holds in virtue of the contents of the beliefs it relates, holds necessarily, and is thus a “logical” relation, at least broadly construed.<sup>16</sup> This conception of evidential support ties it closely to so-called “logical probability,” a term typically associated with Rudolph Carnap. According to Carnap, one set of beliefs (or sentences, as he would say) make another belief probable in virtue of a relation of “partial entailment” between them. He believed that, in principle, such a relation could be completely *syntactically* characterized, and thus one could, in theory, develop an “inductive logic” exactly analogous to deductive logic, that is, a formal system which would enable one to determine, via *a priori* calculations, the degree to which one set of beliefs make another probable.<sup>17</sup>

Today, many epistemologists find such a project extremely dubious. Susan Haack, for example, says that the idea of inductive logic is “prone to paradox at best, perhaps mythical at worst” a sentiment which, I believe, many theorists share.<sup>18</sup> An alternative view is to hold that the relation of evidential support, rather than being a logical relation which is intrinsic to the content of the beliefs it relates (or capable of syntactic characterization), is a contingent relation, or more accurately, a *nomological* connection among beliefs. Very roughly, the view would be that if a set of beliefs evidentially supports some other belief, this is a result, not of a connection between their contents, but of the nomological structure of the physical world. A set of beliefs evidentially supports another, on this suggestion, when there is a strong tendency or propensity for the latter to be true, given that the former are true, where this tendency or propensity is grounded in physical process and laws operative in the world. William Alston advocates a view of

---

<sup>16</sup>Such a conception seems to be offered in Kyburg (1974), as well as Bonjour (1986), and is ambivalently accepted in Fumerton (1995).

<sup>17</sup>See Carnap (1950) and (1952).

<sup>18</sup>Haack (1993), p. 84. See also van Fraassen (1989), ch. 12, for further criticism of the notion of logical probability.

this sort. Having said that a belief is justified only if it is based on a ground, he goes on to say:

Not every grounded belief will be justified, but only one that has an *adequate* ground. To get at the appropriate criterion of adequacy, let's note that a belief's being justified is a favorable status vis-à-vis the basic aim of believing or, more generally, of cognition, viz., to believe truly rather than falsely. For a ground to be favorable relative to this aim it must be "truth conducive"; it must be sufficiently indicative of the truth of the belief it grounds. In other terms, the ground must be such that the *probability* of the belief's being true, given that ground, is very high. It is an objective probability that is in question here. The world is such that, at least in the kinds of situations in which we typically find ourselves, the ground is a reliable indication of the fact believed.... Suffice it to say that I am thinking in terms of some kind of "tendency" conception of probability, where the lawful structure of the world is such that one state of affairs renders another more or less probable.<sup>19</sup>

If we interpret Alston's talk of grounds in terms of our notion of evidence, and if we interpret his reference to the adequacy of a ground in terms of the present notion of evidential support (i.e., a ground is adequate for a belief if it evidentially supports that belief), then Alston is here expressing a view of evidential support which is very different from the first view discussed above. Evidential support, on that view, holds necessarily among beliefs in virtue of their contents, if it holds at all. On Alston's view, it is determined not by the contents of beliefs (or their syntactic form), but by the *world* and its physical laws.<sup>20</sup>

We have seen that a belief's being inferentially justified by certain other beliefs requires at least that the latter provide genuinely good evidence for the former. This we explained in terms of a primitive notion of an evidential support relation holding among

---

<sup>19</sup>Alston (1989), pp. 231-232.

<sup>20</sup>Other epistemologists have held similar views. Nozick, for example, speaks of the "contingency of the evidential tie", and says, "Whether e is evidence for h depends upon what factual empirical connection holds between e and h" (Nozick 1981, p. 261). He goes on to explain this connection in terms of subjunctive "tracking" relations, and probabilistic "approximations" to such relations. We will revisit this tracking notion in Chapter 3.

beliefs. But as we shall see more fully in the next chapter, the mere fact that a person holds certain beliefs which evidentially support some further belief is not sufficient to render that further belief inferentially justified. Consider, for example, an individual, call him Jones, who has followed an involved and convoluted criminal trial in which near conclusive evidence of the defendant's guilt has been presented, and that Jones is convinced of his guilt. Suppose, however, that while Jones believes each of the facts presented in evidence, as a group they are highly complicated and their relevance to the proper verdict is too subtle for him to perceive. Instead, he believes the defendant is guilty simply because she "looks guilty" to him; he thinks her anxious countenance betrays her culpability. It seems, then, that Jones in some sense *has* evidence for his belief that the defendant is guilty, he holds beliefs which in fact evidentially support this conclusion, but his belief is not inferentially justified *by* that evidence.

Certain further conditions of inferential justification have been proposed in the epistemological literature to explain why a belief in such a case is not inferentially justified by the evidence in the subject's possession. One of these conditions, which I shall call the *Basis Condition*, requires that the subject not simply *have* evidence for her belief, but that she *use* the evidence in arriving at her conclusion, that she hold the belief on the *basis* of her evidence. This condition can help to explain why Jones's belief in the defendant's guilt is not justified since Jones's inability to follow the prosecutor's argument has prevented him from using the evidence in his possession to reach his conclusion. A second condition which has been proposed to explain the lack of justification in cases like the present one may be called the *Access Condition*. It requires that a subject be able to "access" his evidence on reflection, and appreciate *that* it supports the belief inferred from it. This condition would appear also to be violated in Jones's case, since Jones is in no position to appreciate the support which his evidence gives to his belief that the defendant is guilty.

Obviously, these two conditions are closely related, and instances in which one is violated will frequently involve the other being violated as well. Still, certain theorists have accepted only one or the other condition, thinking that the two together would make the requirements for inferential justification too stringent. I shall argue in the next chapter that this is a mistake, and that both conditions are indeed necessary for inferential justification. I will then proceed in subsequent chapters to spell out these two conditions in more precise terms, and to respond to certain pressing objections which the resulting account faces.

Before turning to these issues, however, I want to discuss the overarching theoretical orientation which will guide the arguments that follow, and to compare my approach with others that have been taken.

### 3. *The Secure Position.*

It is clear from the appraisals we make of our own beliefs and the beliefs of others that we operate in many everyday situations with a two-fold cognitive desideratum. Specifically, we desire to hold beliefs on the basis of adequate evidence, and to be able to recall or access that evidence on reflection, as the situation demands. We might think of this as a desire to be in a *secure position* with respect to the truth, to have some subjective *assurance* of the truth of our beliefs which we can communicate to others should the need arise. I regard the theory of inferential justification developed in subsequent chapters as an attempt to philosophically explicate this intuitive idea of being in a secure position with respect to the truth.

We can see the desire to be in a secure position manifested in a wide variety of contexts. For instance, it is evident in deliberative situations in which decisions or recommendations must be made which are likely to be challenged by others. Thus, if Smith is part of a committee charged with deciding who should succeed Jones as

president of the college, she will presumably be subject to an expectation that she not make an arbitrary or biased ranking of candidates, that she not reach a decision by flipping a coin or throwing darts at a list of the candidates's names, and so on. Smith's colleagues will expect her to have reasons why she thinks a particular candidate is better than some other, and to be able to produce these reasons if requested to do so (and she will no doubt expect the same of her colleagues). If Smith desires to meet these expectations, she will want to be in a secure position with respect to the truth of her beliefs concerning the relative qualifications of the candidates, that is, she will want to have evidence for these beliefs which she can access or recall on reflection.

The desire to be in a secure position is also present in contexts of directed inquiry. The process of pursuing an inquiry, of raising questions and deliberately searching for answers, seems to require that one keep track of the assumptions, methods, and inferences whereby one conducts that inquiry, that one have a sense of where one stands in one's efforts to gain answers to the questions raised. To take a simple example, a car mechanic might be asked to investigate why a certain car lacks the horsepower to climb only moderately steep hills. Solving the problem, no doubt, will involve checking various possibilities, e.g., whether the engine is properly timed, whether the idle speed is set properly, whether the fuel filter is clogged, and so on. In investigating each of these possibilities the mechanic will have to employ various procedures, from inspecting engine parts with the naked eye, to using various specialized instruments (a timing gun, a tachometer, a diagnostics computer, etc.), and she will have to keep track of what each of these procedures reveals in order to determine what is the correct explanation of the car's poor performance. And no doubt she must be able in the end to explain to the car's owner the nature of the problem, and how she proposes to fix it. Thus, if the inquiry proceeds as it ought, the mechanic will in the end have adequate evidence for an

explanation of the car's poor performance which she can recall or access on reflection. In other words, she will be in a secure position with respect to the truth.

So, in one's efforts at successful deliberation and inquiry, one ideally wants to be in a secure position with respect to the truth. Given that it is a fundamental part of our nature to be deliberative and inquisitive, we have an interest in philosophically scrutinizing this desideratum, and in formulating precise conditions which must be satisfied in order to attain it. As I have indicated, my purpose in exploring the nature of inferential justification is to contribute to an explication of this secure position ideal.

This approach to justification, treating it as a matter of being in a secure position with respect to the truth, differs from the two dominant orientations in epistemology, often called the *traditional* and *naturalistic* approaches, although it has important elements in common with both. I shall describe these, and their relations to my view, in turn.

Traditional epistemology is concerned most centrally with the *critical* task of evaluating our everyday beliefs about ourselves and the world we live in and determining to what extent these are supported by evidence which is free from empirical presupposition. As Laurence Bonjour, a forceful advocate of the traditional approach explains it, "one central aim of traditional epistemology is (a) to decide whether or not we have good reasons for thinking that our various beliefs about the world are true; and (b) if the answer to this first question is affirmative, to say what those reasons are and to explain *why* they are good ones." He goes on to say that this is "*the* central issue, the common thread that connects the concerns of Descartes, Locke, Hume, Reid, and others with twentieth-century epistemology."<sup>21</sup> The philosophers Bonjour cites, particularly and especially Descartes, can be seen as primarily focused on the problem of answering

---

<sup>21</sup>BonJour (1994), p. 284. See BonJour (1985), ch.1, (1992), and (1995) for further discussions of his conception of traditional epistemology. Richard Fumerton is another prominent traditionalist. See his (1985), (1988), and (1995).

*skepticism* concerning our knowledge of the external world, of determining to what extent we are able to respond to one who would challenge us to produce noncircular reasons for our empirical beliefs (i.e., reasons which do not themselves depend on empirical information). Thus, in saying that the business of traditional epistemology is to explain why we have good reasons for thinking that our beliefs about the world are true, BonJour means that the task is to explain this to the satisfaction of a skeptic, that is, to explain this in a way which does not presuppose that any of these reasons *are* in fact good ones.

In addition to holding that it is the business of epistemologists to respond to skepticism concerning our knowledge of the external world, traditionalists quite naturally tend to take the further view that a given individual's *having* knowledge or justified beliefs requires that she herself be in a position to respond to skepticism, that she herself be able to explain to the satisfaction of a skeptic why she has good reasons to think her beliefs are true.<sup>22</sup> But I offer a very different understanding of epistemic notions. Having a justified belief in the sense of being in a secure position with respect to the truth does not require that one be able to respond to skepticism concerning our knowledge of the external world, for that is not something which ordinary people strive to achieve in typical contexts of cognition. In the sorts of contexts in which human cognition normally occurs, subjects are evaluated with respect to whether they are able to access or recall evidence for their beliefs which is adequate to satisfy challenges and requests for evidence that are driven by ordinary human concerns. Think again of the auto mechanic explaining the cause of the car's performance falloff to its owner, or of Smith explaining to her colleagues the reasoning behind her ranking of the candidates for president of the college. These individuals are not responding to skeptical challenges, but to the sorts of requests for evidence and questioning of claims which all of us encounter during the course of everyday affairs.

---

<sup>22</sup>See, for example, BonJour (1985), pp. 8-11, and Fumerton (1988), p. 457.

Despite the clear difference between our respective views, there is an important connection between my secure position approach to justification and the traditional anti-skeptical approach. For we both maintain that having a justified belief requires the subject to be aware, or to be capable of becoming aware, of the evidence which supports that belief. On the traditionalists view, that evidence must be sufficient to satisfy a skeptic concerning our knowledge of the external world, while on my view one need only have the kind and amount of evidence required to get along in one's everyday activities of deliberation and inquiry. But the idea of having a reflective awareness of, or *access* to, the evidence which supports one's beliefs is the same. To use a term frequently encountered in the epistemological literature, we each adopt an *internalist* conception of justification. Our views share the common principle that a subject has a justified belief only if at least some of the conditions or factors responsible for her belief's being justified are accessible or available to her on reflection, only if, that is, she can "tell on reflection" that these conditions or factors obtain.<sup>23</sup> Our position on this matter is opposed by so-called *externalist* views which do not require any conditions or factors of justification to be reflectively accessible in this way. I shall come to them in a moment.

So, my approach shares with the traditionalists the idea that justification must be understood in internalist terms, that a subject must be able to tell on reflection that certain conditions or factors of justification obtain, specifically, that her beliefs are supported by adequate evidence. I disagree with the traditionalists with regard to the kind and amount of evidence which is required to be reflectively accessible to the subject, since they require access to a kind or amount of evidence which is out of proportion to the needs and concerns of ordinary human beings. I focus instead, as I have said, on the idea of being in a secure position with respect to the truth.

---

<sup>23</sup>In truth, as will emerge, I accept this principle only in regard to *inferential* justification. We may ignore this complication for the time being, however.

It is important to distinguish the claim I am making here from two other claims which might easily be confused with it. First, in saying that my treatment of epistemic notions in terms of being in a secure position is advantageous because it ties such notions to the needs and concerns of ordinary human beings, I am not saying that the secure position approach must be correct because it captures the ordinary meanings of terms like 'knowledge' and 'justified belief,' nor am I saying that the traditional approach in epistemology is incorrect or misguided because it misconstrues the ordinary meaning of these terms. My claim is merely that we have a considerable philosophical interest in explicating a desideratum which guides our cognitions in many everyday contexts, i.e., being in a secure position with respect to the truth. Because this is clearly an epistemic notion, it seems right to say that a subject who is in a secure position with respect to the truth in holding a certain belief is in an obvious sense *justified* in holding it. Whether this is the ordinary meaning of 'justified' is not of primary concern. What matters is the explication of the secure position notion itself. Accordingly, the traditionalist approach in epistemology is not false or misguided because it fails to explicate epistemic notions which are of interest in everyday affairs, or which are expressed in ordinary uses of epistemic terms. Rather, the traditional approach is *incomplete*; it does not exhaust the field of epistemological investigation, for there is an important project (important because tied to our everyday concerns) which it does not take up, specifically, the explication of what it means to be in a secure position with respect to the truth.

Second, I do not mean to claim that the traditionalists efforts to respond to skepticism concerning knowledge of the external world are somehow absurd or misguided because they are out of touch with the everyday concerns of ordinary human beings. On the contrary, the endeavor to respond to skepticism is a central part of epistemology, and is a natural problem to focus on in a philosophical investigation of knowledge and justified belief. But again, it is not the only epistemological problem. For

if in ordinary contexts we do not typically evaluate one another with regard to whether we are in a position to respond to skepticism, but with regard to whether we are in a secure position with respect to the truth, then it would seem also to be the business of philosophers to focus on this latter notion and to spell it out in precise terms.

The difference between my approach and that of the traditionalists can be succinctly described in this way: I am interested in philosophically explicating a cognitive desideratum which is apparent when we examine human cognition in some of its natural settings, i.e., in everyday contexts of deliberation and inquiry. Traditionalists, on the other hand, are interested in a problem which arises in *philosophical* contexts, namely, the problem of showing to the satisfaction of a skeptic that we have knowledge of the external world, and they take the having of knowledge to consist in being able to solve this problem.

The fact that I am interested in human cognition as it occurs in its natural setting means that my approach has something in common with the theoretical orientation which currently rivals traditionalism is epistemology, so-called “naturalized epistemology.”<sup>24</sup> The central tenet of this relatively recent movement is that epistemologists should aim not to answer skepticism, but rather to give the best scientific account of the features of ourselves and the world in virtue of which we are for the most part able to depict it accurately in our beliefs.<sup>25</sup> The naturalists thus think of knowledge and justified belief as natural phenomena: they regard a creature’s knowledge or justified beliefs as a natural

---

<sup>24</sup>The term ‘epistemology naturalized’ was introduced in Quine (1969), although the distinction between the traditional approach and that which would come to be called the ‘naturalistic’ approach was recognized in Russell (1940), pp. 10-13. For contemporary examples of the naturalistic approach in epistemology, see Goldman (1986), Kornblith (1993), (1994a), (1994b), and (1995), Kitcher (1992), and Kim (1988).

<sup>25</sup>As Kornblith puts it: “The question of how knowledge is possible... is an empirical question: it is a question of how creatures such as we (given what our best scientific theories tell us we are like) may come to have accurate beliefs about the world (given what our best current scientific theories tell us the world is like)” (Kornblith 1992, p. 298). See also Kornblith (1993), pp. 2-4.

result of its information processing abilities which, of course, are susceptible to fruitful scientific study in the same way that a creature's digestive or circulatory system may be scientifically investigated.

While my approach shares with the naturalists an interest human cognition in its natural setting, the empirical focus of the naturalists typically leads them to accounts of epistemic notions which contrast markedly with both the traditionalist account, and with the proposal I will ultimately recommend. For, in thinking of justified belief as a natural result of a creature's information processing abilities, the naturalists tend to emphasize the ways in which human cognition is similar to the cognition of other species in the animal kingdom. Indeed, some naturalists explicitly state that their goal is to provide an account of epistemic terms which shows that other cognizing animals have justified beliefs (and knowledge) in exactly the same sense as human beings ("human pathways to belief," as Barbara Winters puts it, "are not different in kind from the rest of the animal kingdom"<sup>26</sup>). Since we have in common with our fellow cognizing creatures the fact that we form beliefs<sup>27</sup> about ourselves and our environments through generally accurate and reliable information processing, naturalists thus tend to think of justified beliefs as ones which result from sufficiently reliable cognitive mechanisms, i.e., mechanisms which exhibit a strong tendency to produce true beliefs (the term 'mechanism' here covers "hard-wired" psychological processes, such as are involved in visual perception or memory, as well as learned heuristics, algorithms, and methodologies).<sup>28</sup> This

---

<sup>26</sup>Winters (1983), pp. 217-218. Some of Winters's views will be examined more fully in Chapter 3.

<sup>27</sup>Let us assume for the moment that these other animals have beliefs in the same sense that we do.

<sup>28</sup>A belief would then count as an item of knowledge on this view if it were true, and produced by a sufficiently reliable cognitive mechanism (assuming we ignore Gettier-type worries). It should be acknowledged that this is the crudest sort of naturalistic account of knowledge and justified belief, and it is difficult to find any actual adherents to it in the naturalistic literature. Perhaps the closest is Goldman (1979), but his views have since undergone considerable development (see e.g., Goldman 1986, 1988, and 1993).

*reliabilist* conception of justified belief, as it is usually called, is externalist in the sense in which that term was used above. The naturalists take it to be the business of epistemologists to understand the features of the "world-mind interface" which allows for reliable belief production, and hence for justified belief and knowledge, but they typically do not require ordinary cognizers to have any idea of the character of this interface in order to have justified beliefs or knowledge about the world. A belief's being justified does not require, on their view, that the subject "tell on reflection" that it is reliably produced.

Now, reliable belief production is without question something that we value in many everyday settings. We worry, for example, whether a witness giving testimony in a court of law is reliable, and we give eyewitness testimony more weight than mere hearsay in view of the presumed greater reliability of the former. But we should not be led thereby to think that reliability is the most central epistemic notion. Similarly, we should not let the fact that our cognitions have something in common with the those of other species in the animal kingdom lead us to suppose that *what* our cognitions have in common is the *essence* of justified belief. For human beings have a combination of traits not found in other animals. Specifically, we are deliberative, inquisitive, and perhaps most importantly, social beings, who raise questions and challenge one another's beliefs. This *combination* of traits leads us to desire not simply reliably produced beliefs, but beliefs which allow us, in holding them, to be in a *secure position* with respect to the truth. And this means that we desire not only to have *evidence* for our beliefs, but to be able to access or recall that evidence on reflection. An account of justification which

---

Other naturalistic accounts of knowledge include Nozick's (1981) account of knowledge as belief which "tracks" the truth. A more recent naturalistic approach to knowledge is developed in Plantinga (1993a) and (1993b). Both of these latter theorists eschew talk of justification in explaining the nature of knowledge.

emphasizes too much the similarities between human beings and other cognizing creatures will throw no light on our desire to attain a secure position.

To summarize the discussion thus far, my aim in this dissertation is to set out certain necessary conditions of inferential justification as a way of partially explicating our pretheoretic desideratum of being in a secure position with respect to the truth. We have seen that this approach has important elements in common with the two dominant theoretical orientations in epistemology, the naturalistic and traditional approaches. The connection, on the one hand, between my proposal and that of the naturalists is that they, like I, wish to explicate a desideratum which is manifested in ordinary contexts of cognition, one which is apparent when we investigate human cognition as a naturally occurring phenomenon. But our views differ in that the naturalists are led in their investigation of cognition as a natural phenomenon to emphasize the similarities between adult human cognition and the cognitions of less sophisticated beings, and thus regard reliably produced belief as the fundamental epistemic notion. I, however, am interested in a notion which I think more accurately reflects our everyday epistemic concern, namely, being in a secure position with respect to the truth. On the other hand, my approach shares with traditionalists the idea that in order to have justified beliefs, a subject must be “in touch” with her epistemic situation, she must have some awareness of where she stands in relation to the truth. More technically, both I and the traditionalists adopt an *internalist* conception of justification, meaning that we require a subject to “tell on reflection” that certain conditions or factors of justification obtain. The difference between our views lies in the fact that the traditionalists require a subject to be in touch with her epistemic situation in way which enables her to respond to skepticism concerning our knowledge of the external world, but on my view the subject need only be able to respond to challenges and questions that are driven by ordinary human concerns.

Having sketched my overall approach and its relations to the naturalistic and traditional positions in epistemology, I should briefly consider two objections which adherents of these alternative approaches might press against the one I advocate. The thrust of these objections is that the desideratum of being in a secure position with respect to the truth is philosophically less fundamental or interesting or worthy of attention than the epistemic notions on which the naturalists and traditionalists respectively focus.

First, a naturalist might argue that what is ultimately at issue in our everyday epistemic evaluations is the *truth* or *falsehood* of the beliefs evaluated, and that we are only interested in being in a secure position to the extent that being in such position substantially increases the probability of our beliefs being true: it is only because being in a secure position with respect to the truth is a *reliable* way of coming to have beliefs that we care about it at all. The suggestion is thus that reliability still proves to be the most fundamental notion when it comes to making sense of our everyday epistemic evaluations.

One suggestion which lies behind this argument, that being in a secure position is not the *fundamental* goal of human cognition must, I think, be conceded. For our desire to be in a secure position seems ultimately to serve a more basic desire to have true beliefs and to avoid having false ones. But the argument is incorrect if the further suggestion is intended that because the ultimate goal of cognition is to have true beliefs and to avoid having false ones, the fundamental mode of epistemic evaluation is the evaluation of beliefs with respect to the reliability of their producing mechanisms. If the ultimate cognitive goal *is* to have true beliefs and to avoid having of false beliefs, then *mere* reliability is not our primary concern in epistemic evaluation, because beliefs which are *merely* reliably produced may not be conducive to the goal of truth in the right way. Given our goal of having true beliefs, we desire, in fact, not simply beliefs which are objectively conducive to this goal, but beliefs which we can *recognize* to be conducive to

the goal of truth. For it is only such beliefs that we can trust as being genuinely conducive to that goal. Intuitively, if one has reliably produced beliefs, but one is unable to recognize that they are so produced (or that they are likely to be true for any reason), then one is in a very *insecure* position. In such a position, one has no reasons which one could produce for one's beliefs, and hence would be unable to support them if challenged. In that case, one is left wondering *why* one even holds the beliefs in question. Clearly, one holds the beliefs because one has *faith* that they are reliably produced or likely to be true. But, *ex hypothesi*, one has no evidence for this supposition either. Thus, one is in the position of having simply to blindly trust one's beliefs, which is not a desirable position to be in given one's goal of having true beliefs. So, even if being in a secure position with respect to the truth is not the fundamental goal of cognition, it still seems to take preeminence over mere reliably produced belief in the scheme of everyday epistemic evaluation.

Just as naturalists might argue that being in a secure position with respect to the truth is less important than reliability in understanding epistemic notions, so too might traditionalists argue that the secure position desideratum is from a philosophical point of view less fundamental or interesting than the desideratum of being in a position to respond to skepticism concerning our knowledge of the external world. For the approach I have recommended in this chapter seems vulnerable to the same sorts of arguments which certain traditionalists have directed against externalist and naturalist accounts of epistemic terms, namely, that these accounts simply "change the subject" and focus on epistemic notions which would not have been of interest to the vast majority of philosophers in the mainstream of epistemological tradition. This is stated most bluntly by Richard Fumerton, who, having acknowledged that epistemic terms have a wide variety of uses and meanings, claims that "as philosophers, however, we can and should try to focus on the philosophically relevant uses of epistemic terms. And the

philosophically relevant epistemic concepts are those, satisfaction of which, resolves philosophical curiosity and doubt.”<sup>29</sup> On his view, “the externalist has failed to analyze a philosophically interesting concept of justification or knowledge.”<sup>30</sup> because knowledge and justified belief, as the externalist understands these, are not the sort of knowledge and justified belief which “the philosopher wants.”<sup>31</sup> His suggestion is, apparently, that the only philosophically interesting concept of justification or knowledge is the concept of that kind of justification or knowledge which we seek in a *philosophical inquiry*. And since in philosophically inquiring into the nature of knowledge, epistemologists have traditionally been interested in responding to skepticism concerning our knowledge of the external world, Fumerton is led to regard the concept of knowledge advocated by the traditional epistemologist, i.e., the concept of being in a position to respond to skepticism, as the only philosophically interesting epistemic notion.

My only response is that this seems to me to be too limited a view of the range of philosophically interesting epistemological questions. If it is part of epistemology’s task to examine and to try to make sense out of pervasive features of our cognitive situation, then our striving to be in a secure position with respect to the truth is one such feature which demands philosophical attention.

I will close this section with a brief but important clarification. In discussing the notion of being in a secure position with respect to the truth, I have spoken of this as involving an ability on the part of subjects to respond to challenges or requests for evidence which are driven by ordinary human concerns. This suggests that there is a social dimension to being in a secure position, and that there is in turn a social dimension to justified belief, that having a justified belief is to be defined in terms of being able to

---

<sup>29</sup>Fumerton (1988), p. 457.

<sup>30</sup>Fumerton (1995), p. 180.

<sup>31</sup>Fumerton (1988), p. 457.

respond to such challenges or requests for evidence. I want to emphasize that all of this talk of responding to challenges and defending one's beliefs is intended as a heuristic device to help *locate* the epistemic notion on which I want to focus, and to suggest the drift of its explication. In setting out conditions of inferential justification in the chapters that follow, no explicit reference will be made to such social notions as responding to challenges or satisfying objectors. Being in a secure position with respect to the truth involves holding beliefs on the basis of adequate evidence which one has access to or awareness of on reflection. I take it that one of the primary reasons why we *care* to be in a secure position with respect to the truth is that we are social beings who from time to time call upon one another to provide evidence for our beliefs. But this social explanation of why we desire to be in a secure position will not be explicitly reflected in the account of inferential justification I will develop.<sup>32</sup>

#### 4. *Summary and Overview.*

In this first chapter, I described the project of this dissertation, which is to explore certain necessary conditions of inferential justification as a way spelling out some of the elements involved in being in a secure position with respect to the truth. I suggested that a belief is justified in this way only if the subject holds certain other beliefs which provide adequate evidential support for it in some objectively truth-linked sense. In subsequent chapters, I will try to determine some of the further factors which are required for a belief to be inferentially justified, using the intuitive notion of being in a secure position with respect to the truth as my guide. The discussion will take the following course:

In the next chapter I will examine two conditions of inferential justification, similar versions of which have been widely discussed in the epistemological literature.

---

<sup>32</sup>Thus, in offering an account of inferential justification in terms of being in a secure position with respect to the truth, I am not adopting a so-called "contextualist" approach to justification. See Annis (1978) for a sketch of a contextualist view.

Each will be given a rough initial formulation, and will then be honed into more precise and acceptable shape in succeeding chapters. The first of these conditions, which I call the *Basis Condition*, requires in its initial formulation that a subject's belief be correctly inferentially based on its evidentiary beliefs if the latter are to confer justification on it; the condition requires an actual, psychologically instantiated inferential (and thus causal) relationship between a belief and its evidence. The second condition which I will discuss, called the *Access Condition*, arises from the fact that, as I indicated in the previous section, inferential justification must be understood in so-called *internalist* terms if we are to adequately capture what it means to be in a secure position with respect to the truth. Having an inferentially justified belief requires that the subject be able to tell on reflection that certain conditions or factors of justification obtain. Specifically, the initial formulation of the Access Condition requires a subject to be able to come to know on reflection that her belief is correctly inferentially based on its evidence.

In Chapter 3, I explore the factors involved in a belief's being correctly inferentially based on its evidentiary beliefs, as required by the Basis Condition. A useful first approximation is that a belief is correctly inferentially based on its evidentiary beliefs just in case the latter *evidentially support* and *causally sustain* it. But these causal and evidential relations prove not to be sufficient for correct inferential basing, as demonstrated by various counter-examples. I argue that correct inferential basing requires that a belief be produced by a general inferential disposition, triggered by its evidentiary beliefs, and that this inferential disposition must correspond to or track a relation of evidential support. Chapter 3 closes with an account of the relevant notion of tracking or correspondence.

In Chapter 4, I argue that, on pain of a vicious regress, only justificatory factors which can be noninferentially known to obtain may be required to be reflectively accessible to the subject in the proper formulation of the Access Condition. Since one

cannot, for reasons which will emerge, in general noninferentially know whether one's belief is correctly inferentially based on its evidence, the initial formulation of the Access Condition given in Chapter 2 must be substantially reformulated. Specifically, the Access Condition must in the end require only that the most salient or crucial *necessary conditions* of a belief's being correctly inferentially based on its evidence be reflectively accessible. I argue that, at a given moment, the most salient factors responsible for a belief's being inferentially justified are the relations of causal sustenance and evidential support between it and its evidentiary beliefs, and that these relations are both reflectively accessible. The Access Condition in its final formulation thus requires merely that a subject be able to tell on reflection that her belief is evidentially supported and causally sustained by its evidentiary beliefs.

In Chapter 5, I respond to objections to the account of inferential justification offered in the earlier chapters. The objections fall into two categories, the first alleging that the conditions of the account are not satisfiable in principle, the second that the conditions are not satisfied in practice. The thrust of the first objections is that neither the relation of evidential support nor the relation of causal sustenance is the *sort of thing* which could be reflectively accessed. I argue that it is only on a mistaken "Cartesian" picture of reflection that these relations are not accessible, and that if one adopts instead a "naturalistic" account of reflective access, then there is no difficulty in supposing that we do on occasion have access to the required relations. I then consider the obvious objection that internalist theories of justification are committed to the Cartesian view of reflection, and argue that internalist accounts can still make sense of what it means to be in a secure position with respect to the truth if the switch is made to a naturalistic conception of reflective access.

The second objections maintain that due to ordinary human deficiencies in information storage capacity and reflective power, very few of a typical adult's beliefs

satisfy either the Basis or the Access Conditions. If these conditions were indeed necessary for justification, then most people would have very few (if any) inferentially justified beliefs, and consequently very few beliefs which are justified by any means, since a great many of an individual's beliefs must be inferentially justified if justified at all. But, the argument concludes, the correct standards of justification must be easy enough to satisfy that most ordinary people have a substantial number of justified beliefs. I have a two-fold response to these objections. First, the extent to which ordinary people fail to satisfy my two conditions is exaggerated. Second, we must keep in mind that the stringency of the epistemic standards we intuitively apply may not be constant, but may fluctuate depending on circumstances. It could be that there are situations in which we intuitively feel that some individual has an inferentially justified belief despite her failure to satisfy the Basis and Access Conditions because we think it appropriate, for what ever reason, to apply a less stringent epistemic standard to her case. It is nevertheless clear that there are many circumstances in which we place a high premium on being in a secure position with respect to the truth, and in at least these circumstances my conditions ought to be applied, regardless of our apparent cognitive limitations.

## Two

### Outline of a Theory

As I described in the first chapter, my aim in this dissertation is to explore and refine certain necessary conditions of inferential justification, thereby spelling out in precise terms some of what it means to be in a secure position with respect to the truth. My purpose in this chapter is to give an initial sketch of the two conditions which will be my focus, and to explain how they relate to the secure position desideratum.

The first section of the present chapter is devoted to what I call the *Basis Condition* of inferential justification. This condition requires of an inferentially justified belief that it be correctly inferentially *based on* its evidence, which means in part that it be *causally sustained* by that evidence. I shall try to show why a “causal condition” of this sort is necessary in order to be in a secure position with respect to the truth, and I shall compare the specific form which my Basis Condition takes with similar causal conditions that have been discussed in the literature.

In the second section I discuss the so-called *Access Condition* which, together with the Basis Condition, explicates the core of what it is to be in a secure position with respect to the truth. This condition requires in its initial formulation that a subject be able to “know on reflection” that her belief is correctly inferentially based on its evidence if it is to be inferentially justified. This reflects the fact that the notion of a secure position must be understood in *internalist* terms, that certain conditions or factors of justification must be reflectively accessible to a subject if she is to be in a secure position with respect to the truth. The initial requirement given here that a subject be able to know on reflection that her belief is correctly inferentially based on its evidence will serve as a place holder

until a more detailed examination of the reflective accessibility of justificatory factors can be provided in Chapter 4.

In the third section of this chapter, I will briefly touch on some details which, in addition to the two conditions I will discuss, would have to be filled in to give a complete account of inferential justification. Since these matters fall outside the scope of this dissertation, my aim will not be to offer specific proposals, but merely to identify certain issues in the neighborhood of the topics on which I am focusing, and to describe avenues of future research.

### *1. The Basis Condition.*

As we saw in the first chapter, if an inferentially justified belief is one which puts the subject in a secure position with respect to the truth, then such justification requires that the subject have in her possession adequate evidence for the truth of her belief. Thus, we have the following minimal condition on inferential justification:

(IJ1) S's belief that  $p$  is inferentially justified by her beliefs  $B_1$  through  $B_n$  only if her beliefs  $B_1$  through  $B_n$  evidentially support her belief that  $p$ .

Our problem is to determine some of the further details which would have to be added to (IJ1) to obtain a fuller understanding of inferential justification. We might begin by noting two matters on which (IJ1) is silent. First, there is no mention of whether S's beliefs  $B_1$  through  $B_n$  are themselves epistemically justified, but presumably they must be if they are to confer justification on the belief that  $p$ . Second, (IJ1) makes no mention of the potential role of "defeating evidence" in determining the epistemic status of the belief that  $p$ . To illustrate, suppose S believes (justifiedly) that Jones is a member of the National Rifle Association, and that 90% of the members of the National Rifle Association would classify themselves as politically conservative. It would seem that these beliefs evidentially support the further belief that Jones would classify himself as

politically conservative in the manner required by (IJ1). But suppose in fact that S also believes (once again, justifiedly) that Jones has voted for every Democratic candidate for president since George McGovern. Then it seems that S is *not* inferentially justified in believing that Jones would classify himself as politically conservative. Her evidence concerning Jones voting behavior defeats her justification for that belief.<sup>1</sup> (IJ1) has no way of accounting for this as it stands.

I mention these issues only to indicate that I will ignore them in what follows. While I shall briefly return to them later in the chapter, I am interested mainly in two further respects in which (IJ1) requires supplementing. These concern the relations which must exist between S, her belief that p, and her beliefs B<sub>1</sub> through B<sub>n</sub>, in addition to the relation of evidential support specified in (IJ1), if the latter beliefs are to confer justification on the former belief. In this section, I will focus on the first of these modifications to (IJ1), namely, that there must be a certain causal connection between S's belief that p and her beliefs B<sub>1</sub> through B<sub>n</sub>. The second required addition to (IJ1), that S must be able to tell on reflection that her belief that p is related to its evidence in the appropriate way, will be discussed in the next section.

The requirement that there must be a certain sort of causal relation between a belief and its evidence develops out of the observation that if S is to truly be in a secure position in holding her belief that p, her beliefs B<sub>1</sub> through B<sub>n</sub> must not only evidentially support her belief that p, but S must, as it were, "see the connection" between these beliefs, and believe that p for that reason. As Hilary Kornblith suggested in his now widely discussed refutation of the "arguments on paper thesis," it is not enough that a person have in her possession an "argument" for the belief that p, in the form of further beliefs which adequately support it; inferential justification requires a psychologically

---

<sup>1</sup>See Pollock (1986) and (1990) for detailed discussions of the nature and classification of various types of defeating evidence. Pollock's distinction, familiar to some readers, between "undercutting defeaters" and "rebutting defeaters" will not concern us here.

realized inference process, i.e., a genuine causal connection, between a belief and its evidence.

The following case, due to Kornblith, illustrates this point. A certain individual, Alfred, believes both that *p* and that if *p*, then *q*, and also believes that *q*. While Alfred's belief that *q* is evidentially supported by the former beliefs in the manner required by (IJ1), he does not believe that *q* *because* it is evidentially supported by these beliefs. Rather, he believes it because he likes the sound of a sentence expressing it.<sup>2</sup> (As Kornblith tells it, Alfred distrusts *modus ponens*, and thus is not inclined to infer the belief that *q* from his beliefs that *p* and that if *p*, then *q*, but we could equally imagine that in forming his belief that *q*, he simply fails to recall these other beliefs, and hence they play no role in the formation of the belief that *q*.)

Kornblith rightly suggests that Alfred's belief that *q* is not justified because of the *way* in which he comes to hold it, despite his having other beliefs which provide evidential support for it. As he puts it, "Alfred is justified in believing that *q* only if his belief that *q* *depends* on his beliefs that *p*, and if *p*, then *q*." He goes on to say that the "notion of belief dependence must be accounted for by looking at the belief states of persons and, in particular, the relations between them. Questions about the justification of beliefs are thus intimately tied to questions about the sorts of processes responsible for the presence of those beliefs."<sup>3</sup>

Let us focus for a moment on Kornblith's claim that a belief must *depend on* certain other beliefs if it is to be justified by them. This means, first of all, that the beliefs which confer justification must be crucially involved in the causal history of the belief on which that justification is conferred. More simply, if a belief that *p* is inferentially justified by a set of beliefs  $B_1$  through  $B_n$ , then the latter must be the cause of the former,

---

<sup>2</sup>See Kornblith (1980).

<sup>3</sup>Ibid., p. 602.

or, alternatively, they must explain the existence of the former. Of course, this is not to say that the beliefs  $B_1$  through  $B_n$  must be (logically or physically) sufficient for the existence of the belief that  $p$ . These beliefs will in the typical case be part of a complex of factors responsible for the belief that  $p$ , but they will be, perhaps, readily identified as the most salient or prominent factors. For example, while at a furniture store I might infer the belief that a certain sofa will not fit through my doorway from beliefs about the dimensions of the sofa and about the width of my doorway. The latter beliefs might be identified as the cause of the former belief, even though other factors are clearly involved as well, such as, for instance, the fact that I am *interested* in whether the sofa fits through my doorway.

As Kornblith acknowledges, inferential justification involves more than a mere causal connection between a belief and its evidence. What is required is not just that the belief that  $p$  causally depend on the beliefs  $B_1$  through  $B_n$ , but that it causally depend on them in virtue of an *inference* which the subject makes. For example, I might believe that someone is in my house because I hear noises within. This belief might cause me to open the door, which in turn might cause me to believe that no one is in the house.<sup>4</sup> While my belief that no one is in the house causally depends on my belief that someone is in the house, it does not depend on that belief in the manner which is necessary for inferential justification, since I could not be said to have inferred that no one is in the house from my belief that someone is in the house.

As a way of intuitively specifying the kind of dependency relation involved in inferential justification, some theorists maintain that a belief must be held *on the basis of*, or *in light of* its evidence.<sup>5</sup> Varying this terminology slightly, I will say that inferential

---

<sup>4</sup>This example is borrowed from Armstrong (1968), p. 195. He in turn borrowed it from G. E. Moore (see Moore 1962, p. 7).

<sup>5</sup>Such phrases are used, for example by Alston (1989), chap. 9, and Swain (1979) and (1981).

justification requires that a belief be held on the *inferential* basis of its evidentiary beliefs, to distinguish the relevant basing relation from the sort of relation that might be involved when a belief is based, say, on a nonpropositional impression or experience.<sup>6</sup> What exactly inferential basing involves, over and above a simple causal connection between a belief and its evidence, is difficult to specify. Alston suggests that a belief is based on a ground (his term for evidence) when “the belief forming ‘process’ or ‘mechanism’ is *taking account* of that ground or features thereof, being *guided* by it.”<sup>7</sup> For the time being, such descriptions (i.e., this talk of a belief’s being formed by a process which takes into account, or is guided by, certain other beliefs) will have to suffice for conveying the general idea of inferential basing. In the next chapter I will delve into the matter in much more detail.

So, we have reached a preliminary conclusion that (IJ1) must at least be supplemented by adding a condition that the subject’s belief be *inferentially based* on its evidence. This suggests:

- (IJ2) S’s belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if
- (i) her beliefs B<sub>1</sub> through B<sub>n</sub> evidentially support her belief that p, and,
  - (ii) her belief that p is inferentially based on her beliefs B<sub>1</sub> through B<sub>n</sub>.

The difficulty with this suggestion, however, is that it simply tacks on to (IJ1) a requirement that the belief that p be inferentially based on the beliefs B<sub>1</sub> through B<sub>n</sub>, without integrating these two conditions in any way. To illustrate the problem, suppose that S has a belief of the form p ∨ (p&q), and infers from this a belief that p, *not* however via a deductively valid inference process, but simply because she is inclined to infer the first disjunct of disjunctive propositions which she believes.<sup>8</sup> Her belief that p is

---

<sup>6</sup>I leave it as an open question whether the basing relationship is significantly different in these two sorts of cases. But I suspect that the “connecting beliefs” which I shall argue are involved in inferential basing are not required in order for a belief to be based on a nonpropositional experience in a way which is appropriate for justification.

<sup>7</sup>Alston (1989), p. 229.

<sup>8</sup>This example is derived from Swain (1988). We shall return to it in the next chapter.

presumably both inferentially based on, and evidentially supported by, her belief that  $p \vee (p \& q)$ , but it would not be inferentially justified by the latter belief. Her belief that  $p$  is not *correctly* inferentially based on its evidence. Conditions (i) and (ii) in (IJ2) are necessary for a belief to be correctly inferentially based on its evidence, but they are not sufficient. For the time being, then, I suggest replacing conditions (i) and (ii), by a single condition, which I call the *Basis Condition*:

(BC) S's belief that  $p$  is inferentially justified by her beliefs  $B_1$  through  $B_n$  only if her belief that  $p$  is correctly inferentially based on her beliefs  $B_1$  through  $B_n$ .<sup>9</sup>

I shall leave to the next chapter the problem of determining what factors, beyond those specified in (i) and (ii) of (IJ2), are involved in a belief's being correctly inferentially based on its evidence. According to the account I will develop, a belief is correctly inferentially based on certain other beliefs just in case it is causally sustained by them, and the inference process which generated it "corresponds to" or "tracks" a relation of evidential support between it and the beliefs which sustain it. The details of this account do not matter for our present purposes, however.

I turn now to consider some features of my Basis Condition which might lead some theorists to reject it. I shall argue, however, that accounts which reject the Basis Condition are unable to adequately capture the intuitive notion of being in a secure position with respect to the truth.

First, the Basis Condition as I understand it is an instance of what Robert Audi calls a "sustaining requirement on indirect justification."<sup>10</sup> Inferential basing requires

---

<sup>9</sup>In view of the fact that a belief can be justified at one time, and unjustified at another, it would, strictly speaking, be more correct to include explicit time references in the statement of this condition ("S's belief that  $p$  is inferentially justified at a time  $t$  by her beliefs  $B_1$  through  $B_n$ ...", and so on). For the time being, I will continue to suppress such time indices, including them only in the final formulations of this and the other condition to be discussed shortly, the Access Condition.

<sup>10</sup>Audi (1993), p. 215. Audi uses the term 'indirect justification' in place of 'inferential justification.'

that a belief be (at least in part) *sustained* by its evidentiary beliefs, which, all else being equal, means that, if even one of these evidentiary beliefs were no longer held, the belief in question would no longer be held either.<sup>11</sup> If I infer that a certain friend is dishonest, and continue to hold this belief long after I have forgotten the specific instances of dishonest behavior which lead me to this conclusion, then my belief, while *once* inferentially based on evidence, no longer satisfies the Basis Condition, and is in consequence not inferentially justified on my view. Not every theorist is willing to impose a sustaining requirement of this sort on inferential justification. Swain, for example, is willing to allow that an inferentially justified belief may be based on “distant reasons,” i.e., beliefs which once gave rise to that belief, but which are now no longer held, and thus are unable to currently sustain it.<sup>12</sup>

Swain’s suggestion that inferentially justified beliefs may be based on distant reasons might seem more plausible than my Basis Condition in light of obvious limitations in the typical human being’s ability to retain evidence in long term memory. Consider, for example, my belief that Thomas Jefferson founded the Democratic party. If truth be told, I have done a terrible job of keeping track of evidence for this belief, so much so that it is highly likely that it is not sustained by any other of my beliefs; I do not hold it on the inferential basis of any evidence, but simply for the vague reason that I remember having read somewhere that it is true. And, in all honesty, I am in a similar position with respect to a great many of my other beliefs (though I may not be in such a position with respect to the majority of my beliefs). I believe that Albert Einstein published his special theory of relativity in 1905, that the Earth on average is 93 million miles from the Sun, that Istanbul was once named ‘Constantinople.’ and so on. These beliefs are not causally sustained by evidence in the strict sense. They have a kind of

---

<sup>11</sup> I say “all else being equal” in recognition of the possibility of overdetermination and the like.

<sup>12</sup> Swain (1981), pp. 87-89.

inertia, if you will, which allows them to remain stored in long term memory without the aid of any specific evidence.<sup>13</sup> If I were asked why I hold these beliefs, I would be at a loss. My only response would be that I *remember* having read or heard somewhere that they are true. On my interpretation of the Basis Condition, none of these beliefs are inferentially justified since none are currently sustained by evidence. According to Swain, however, these beliefs might well be inferentially justified, if they were produced by appropriate “distant reasons.”

The problem with Swain’s suggestion is that many people, myself included, consider it a *defect*, a cause for some consternation, when they discover that a certain belief is not in fact held on the basis of any evidence. One may be quite confident that one *once* was aware of evidence for the belief, as I am, for example, in the case of my belief that Thomas Jefferson founded the Democratic party. But this, in at least some instances, makes it all the more worrisome that one has now discarded that evidence. One feels that one really *ought* to have done a better job keeping track of the evidence for one’s belief; one would be in a *more* secure position with respect to the truth if one had retained the evidence on which one originally based the belief. A theory of inferential justification ought to capture this sense of insecurity which goes with holding a belief for which one has failed to retain evidence. It seems then that if holding a belief on the inferential basis of certain other beliefs is to contribute to one’s being in a secure position with respect to the truth, this must involve a relation of causal sustenance between one’s evidentiary beliefs and the belief one bases on them.

Some theorists would go even farther than Swain in rejecting my Basis Condition. Keith Lehrer, for example, holds that an inferentially justified belief need not be causally related to its evidence in any way; all that matters for inferential justification

---

<sup>13</sup>Psychologists refer to this as the phenomenon of *belief perseverance*, the tendency of beliefs to be held long after their evidence has been forgotten, or even discredited. I will discuss it further in Chapter 5.

is that the subject in fact have adequate evidence for the belief in her possession. Lehrer claims, somewhat misleadingly, that when “a man’s justification for his belief is based on evidence, then he believes what he does because of the evidence”, but then goes on to say:

This suggests a *causal* account of what is involved when the justification of a belief is based on evidence. The notion of a justification being based on evidence would then be explicated in causal terms. Following this proposal, a man’s justification for his belief is based on certain evidence if and only if his belief is causally related in some specified way to the evidence. How to specify the exact way in which the belief must be causally related to the evidence would remain a problem on this approach, but it would be a problem of detail rather than of principle. All such theories must be rejected however.<sup>14</sup>

Lehrer’s argument for this last statement, that all such causal theories must be rejected, turns on a number of purported counter-examples, the most developed of which involves the case of a gypsy lawyer “with the gypsy’s traditional reverence for reading the cards.”

This lawyer

has a client who appears guilty to the impartial observer. In fact, imagine that eight murders have been committed, all of them alike, carried out with the same sort of weapon under the same circumstances. There is no doubt about the fact that seven of these murders were committed by the client: that is conclusively established, and the eighth is like the other seven in every detail. The crimes are hideous, and it is assumed by all concerned including the lawyer that his client committed the eighth. However, one evening the lawyer is checking with the cards, reading them with gypsy skill and conviction, when the cards tell him that his client is innocent of the eighth murder though guilty of the others. He never doubts the cards. His conviction of innocence leads him to reconsider the evidence. He discovers a valid though complicated line of reasoning from the evidence to the conclusion that his client did not commit the eighth murder.... He freely admits, however, that the evidence... is not what convinced him of the innocence of his client, and, indeed, would not convince him now were he not already convinced by the cards.... The evidence is quite conclusive, as shown by his complicated chain of reasoning, but even he would find himself unable to believe his client could be innocent of the eighth murder were it not for the fact that the cards told him it was so.... His conviction could not be increased by his consideration of the evidence because he was already completely convinced. On the other hand, were his faith in the cards to collapse, then emotional factors which influence others would sway him too. Therefore, the

---

<sup>14</sup>Lehrer (1974), pp. 122-123.

evidence which completely justifies his belief does not explain why he believes as he does, his faith in the cards explains that.... Since his belief is completely explained by his faith in the cards, the evidence is irrelevant to explaining why he believes as he does.<sup>15</sup>

Lehrer insists in this passage that the lawyer's belief in his client's innocence is inferentially justified despite the fact that there is no causal connection between that belief and his evidentiary beliefs. The lawyer's evidence neither produces nor sustains the belief that his client is innocent of the eighth murder. Thus, Lehrer would say that the gypsy lawyer's belief constitutes a counter-example both to my Basis Condition, and to Swain's suggestion that inferentially justified beliefs must at least be based on distant reasons (if not currently sustained by evidence).

The difficulty with Lehrer's argument, as I see it, is that he focuses exclusively on the fact that the lawyer's belief is *supported by evidence*, and quite fails to acknowledge that his belief is nevertheless *subject to criticism*. The lawyer's belief is subject to criticism because it really *ought* to be held on the basis of his evidence. (It's true that Lehrer speaks of the lawyer's *justification* being based on his evidence, but this is not the same as saying that his *belief* is based on his evidence.<sup>16</sup>) Intuitively, the lawyer is not in a secure position with respect to the truth in believing that his client is innocent for two reasons. First, forming beliefs on the basis of the cards is not a very secure way of forming beliefs, since such beliefs are (presumably) not very likely to be true. While the lawyer is in fact eventually led by the cards to a true belief concerning his client's innocence, his belief is true by the happiest of accidents. It is only because the cards happened on this occasion to lead him to believe that his client is innocent that he believes something true. If the cards had indicated rather that his client was guilty, the lawyer would have believe that instead. Second, as the case is described, if the lawyer were led to distrust his cards, he would give up the belief in his client's innocence, and would thus

---

<sup>15</sup>Ibid., pp. 124-125.

<sup>16</sup>Audi (1993), p. 221. makes this point.

give up a belief for which he *has* good evidence. But, intuitively, being in a secure position with respect to the truth should protect one from giving up beliefs which are in fact well supported.

Thus, if we say, following Lehrer, that the lawyer's belief in his client's innocence is inferentially justified, then we do not do justice to the sense that he did not go about forming that belief in the right way. An account of inferential justification ought to reflect the fact that the lawyer's belief is open to criticism due to the insecurity of his position in holding that belief. It seems therefore that the lawyer is *not* inferentially justified in believing his client to be innocent of the eighth murder.

In fairness to Lehrer, it should be acknowledged that he most likely has the *sincere intuition* that the lawyer's belief in his client's innocence is justified. In view of this fact, it might be helpful to follow Roderick Firth in distinguishing between two different types of justification, which he calls propositional and doxastic justification.<sup>17</sup> On Firth's view, one can regard a *proposition* as justified *for* a given subject, if that subject possesses adequate evidence for its truth. But the subject is justified *in believing* that proposition only if that belief is held on the basis of her evidence. Using this distinction, we can say that Lehrer is interested in propositional justification, and we should interpret his insistence that the lawyer's belief in his client's innocence is justified as strictly speaking a claim, not about the lawyer's belief, but about the proposition he believes, to wit, that this proposition is supported by his evidence. We can then consistently acknowledge Lehrer's intuition, but continue to hold that the lawyer is not justified in believing that his client is innocent of the eighth murder because that belief is not correctly inferentially based on the lawyer's evidentiary beliefs. The case of the gypsy

---

<sup>17</sup>Firth (1978). In actuality, Firth uses the terms 'propositional warrant' and 'doxastic warrant'.

lawyer does not therefore constitute a genuine counter-example to my Basis Condition, insofar as it is a condition of *doxastic*, not *propositional*, justification.<sup>18</sup>

Despite the fact that the Basis Condition seems to make sense as part of an account of what it means to be in a secure position with respect to the truth, there is still a major worry concerning this condition to which we shall have to return. For, many theorists would still claim that the Basis Condition is too implausible from the standpoint of ordinary human psychology, given the relative prevalence of beliefs based only on distant reasons. Since I have admitted that beliefs of the sort discussed above, such as my belief that Thomas Jefferson founded the Democratic Party, or that Einstein published his special theory of relativity in 1905, are not strictly speaking inferentially justified, given their failure to satisfy (BC), and since, plausibly, such beliefs must be inferentially justified if they are justified at all, (BC) implies that quite a few beliefs which would intuitively *seem* to be justified are not in fact justified after all. One important question is whether this result undermines (BC)'s claim to be an adequate (partial) explication of what it is to be in a secure position with respect to the truth. I shall argue in Chapter 5 that it does not.

This concludes my preliminary examination of the Basis Condition and the issues surrounding it. I turn now to the second important necessary condition of inferential justification that I want to discuss, the Access Condition.

## 2. *The Access Condition.*

As we discussed in the first chapter, there are many situations in which others expect us to be able to produce evidence or give reasons for our beliefs. Thus, if one is truly to be in a secure position with respect to the truth, one must be able to recall or access the evidence on which one's beliefs are based. This means that a theory of

---

<sup>18</sup>Or at any rate, it is a condition of *inferential* doxastic justification.

inferential justification, to the extent that it aims to explicate the notion of a secure position, must be an *internalist* theory, it must require that certain conditions or factors of justification be accessible to the subject on reflection. In this section I will begin to sketch the internalist aspects of the theory of inferential justification which emerges as we attempt to make precise sense of what it means to be in a secure position with respect to the truth.

I want to begin by pointing out a couple of natural misunderstandings of the distinction between internalist and externalist theories. First, sometimes the terms ‘internalism’ and ‘externalism’ are used in a way which gives the misleading impression that their ranges of application are rigidly demarcated and mutually exclusive. In contrast, I think it is more accurate to view internalism and externalism as admitting of degrees. One can distinguish between purely internalist theories, purely externalist theories, and what might be called *mixed* or *hybrid* theories.<sup>19</sup> Thus, consider this passage by Matthias Steup:

What makes an account of justification internalist is that it imposes a certain condition on those factors that determine whether a belief is justified.... The condition in question requires J-factors<sup>20</sup> to be *internal to the subject's mind* or, to put it differently, *accessible on reflection*. What makes an account of justification externalist, in contrast, is that no such condition is imposed. According to externalism, J-factors need not be internal to the subject's mind or accessible on reflection.<sup>21</sup>

The difficulty with this description of the internalism/externalism distinction is that there would appear to be several “J-factors” relevant to justification in a given case, and it is a substantive question whether all or merely some of these factors must be reflectively accessible. With respect to a particular inferentially justified belief we might initially

---

<sup>19</sup>See Audi (1993), p. 304, for a discussion of this point. See also Alston (1993) and Conee (1988) for examples of such mixed or hybrid theories. The views of these theorists will be examined in Chapter 4.

<sup>20</sup>The phrase ‘J-factors’ is short hand for “factors that determine whether a belief is justified.”

<sup>21</sup>Steup (1996), pp. 84-85.

identify as relevant factors the fact that the subject holds certain evidentiary beliefs, and the fact that her belief is correctly inferentially based on this evidence. We might go on to recognize all of the various factors (yet to be identified) which are involved in a belief's being correctly inferentially based on its evidence (e.g., the fact *that* the evidentiary beliefs provide support for the belief in question, or the fact that they causally sustain it). And finally, if a belief's being justified requires that J-factors for that belief be reflectively accessible, then the fact that a particular J-factor is reflectively accessible constitutes a further J-factor with respect to that belief.

A *pure* internalist theory of inferential justification would require that *all* of the factors responsible for a belief's being inferentially justified be accessible to the subject on reflection; a pure externalist theory would not require any of these factors to be thus accessible; finally, a hybrid internalist-externalist theory would require that some but not all of the J-factors with respect to a given belief be reflectively accessible. Now, we have seen that if a theory of justification is to make precise sense out of what it means to be in a secure position with respect to the truth, then it must require that at least some conditions or factors of justification be reflectively accessible, thus rendering pure externalism unacceptable as account of the nature of such a position. But it is important to realize that pure internalism is also implausible as a way of explaining the nature of a secure position. For pure internalism seems committed to *both* of the following principles:

- (I) If *F* is a J-factor for S's belief that *p*, then S can come to know on reflection that *F* obtains.
- (II) That S can come to know on reflection that *F* obtains is itself a distinct J-factor for S's belief that *p*.

But these principles lead to a vicious regress. For suppose that S has an inferentially justified belief that *p*. Then, by (I), any J-factor *F* responsible for S's belief that *p* being inferentially justified must be such that S can come to know on reflection that *F* obtains.

But, by (II), the fact that S can come to know on reflection that a particular J-factor  $F$  obtains is itself a distinct J-factor  $F'$ . By (I), therefore, S must be able to come to know on reflection that  $F'$  obtains, i.e., S must be able to come to know on reflection that she can come to know on reflection that  $F$  obtains. By (II), the fact that S can come to know on reflection that  $F'$  obtains is a further J-factor  $F''$ , and by (I), S must be able to come to know on reflection that  $F''$  obtains, i.e., she must be able to come to know on reflection that she can come to know on reflection that  $F'$  obtains, i.e., she must be able to come to know on reflection that she can come to know on reflection that she can come to know on reflection that  $F$  obtains, and so on indefinitely. Clearly, at some point in the infinite reiteration of the construction 'S can come to know on reflection that...', we will arrive at a that-clause so convoluted that S will be unable to comprehend it, and thus will be incapable of having the reflective knowledge in question. This shows that either (I) or (II) must be false.<sup>22</sup>

---

<sup>22</sup>One could try to argue that an infinite regress of the sort generated here is not vicious, provided that the subject is in a position to "see" that each proposition in the hierarchy is true, even if she is incapable of simultaneously believing them. To illustrate, consider some proposition which I know to be true—for instance, that Nixon was once President. Relying on the fact that a proposition  $p$  is equivalent to the proposition that it is true that  $p$ , we can generate the following regress: (i) It is true that Nixon was once President; (ii) It is true that it is true that Nixon was once President; (iii) It is true that it is true that it is true that Nixon was once President; and so on. While I am not capable of simultaneously believing all of these propositions, I can *see* that each of them is *true*, given the truth of the proposition that Nixon was once President, thus rendering the regress harmless. (See Sosa 1993, p. 150, for further discussion along these lines.) But this observation is of no help to pure internalism, since the regress which besets this position is dissimilar to the present regress in an important respect. For in the regress which pure internalism generates, we eventually hit a stage after which all subsequent propositions are clearly seen to be *false*, not true. For if I consider the hierarchy of propositions--(i) I can come to know on reflection that  $F$  obtains; (ii) I can come to know on reflection that I can come to know on reflection that  $F$  obtains; (iii) I can come to know on reflection that I can come to know on reflection that I can come to know on reflection that  $F$  obtains; and so on--I eventually reach a point after which I know that all subsequent propositions are false, for they attribute to me reflective capacities which I know I do not possess.

It seems to me that the problematic principle is (I), but in actuality, there is little difference which of these principles is rejected. For if the pure internalist were to try to escape the regress by giving up (II), this would amount to a concession that certain very important facts which contribute to a belief's being inferentially justified, namely, facts about what is reflectively accessible to the subject, need not themselves be reflectively accessible to that subject (since they are not J-factors in the relevant sense). This would constitute a rejection of pure internalism, regardless of whether we call these facts "J-factors."

In any event, a plausible internalist theory of inferential justification must require that some but not all factors responsible for a belief's being inferentially justified must be accessible to the subject on reflection. The only plausible account of what it is to be in a secure position with respect to the truth is thus a *hybrid* internalist-externalist theory which recognizes that certain important justificatory factors need not be reflectively accessible.

This brings us to the second misunderstanding of the internalism/externalism distinction which is occasionally encountered in the literature. It is sometimes said that internalism is the view according to which a subject holds a justified belief only if she can tell on reflection *that* she is justified in holding it; externalism, on the other hand, does not require her to know this. Roderick Chisholm, for example, has said that "the concept of justification is *internal* and *immediate* in that one can find out directly, by *reflection*, what one is justified in believing at any time."<sup>23</sup> And again, "the internalist assumes that, merely by reflecting upon his own conscious state, he can formulate a set of epistemic principles that will enable him to find out, with respect to any possible belief he has, whether he is *justified* in having that belief."<sup>24</sup> I have two problems with this

---

<sup>23</sup>Chisholm (1989), p. 7.

<sup>24</sup>*Ibid.*, p. 76.

interpretation of the internalist position (and the implied interpretation of the externalist position).

First, Chisholm speaks of the internalist formulating certain principles which allow him to sort his beliefs into those which are justified and those which are not. Does this imply that the ordinary person is required to be able to formulate such principles? If so, then a person has justified beliefs only if she is able to formulate on reflection what is in effect a philosophical *theory* of justification which she can apply in determining which of her beliefs is justified. This is clearly too much to expect of an ordinary person. It does not seem right to describe the sort of reflective access which an ordinary person has to her epistemic situation in terms of her being able to formulate a single set of principles which are strong enough to determine whether any given belief she holds is justified or unjustified. Thus, in speaking of “formulating principles” by “reflecting on his conscious state,” Chisholm must mean that the internalist holds a methodological view concerning the manner in which one may go about constructing a theory of justification, namely, that one may do this simply by reflecting on one’s own conscious state. But if an internalist theory is to provide an account of what it means to be in a secure position with respect to the truth, then it cannot be a view merely about what epistemologists are in a position to do in constructing theories of justification. It must lay down requirements which ordinary people are required to satisfy.

Second, if a subject is to determine by reflection *that* she is justified in holding some belief, as Chisholm requires, then she would seem to have to be capable of determining on reflection that this belief has all of the features which are necessary for making a belief justified, including the fact that she can tell on reflection that it is justified. But then the fact that she can tell on reflection that she can tell on reflection that her belief is justified would seem itself to be a feature which is necessary to make her belief justified, implying that she would have to be able to tell on reflection that she can tell on

reflection that she can tell on reflection that her belief is justified. Thus, we seem led into the same regress which besets the pure internalist position.

So, in making sense of what it means to be in a secure position with respect to the truth, we must not require that the subject be able to formulate a comprehensive set of epistemic principles, or that she be able to tell on reflection *that* her belief is inferentially justified, or that all of the factors responsible for her belief's being justified obtain. But then what justificatory factors *should* be required to be reflectively accessible? My provisional suggestion is that we take our cue from the Basis Condition. Since it requires a subject's belief to correct inferentially based on her evidentiary beliefs, the most natural thing would be to require *this* to be reflectively accessible.<sup>25</sup> Accordingly, I will adopt for now the following requirement on inferential justification, which I will call the *Access*

*Condition*:

(AC) S's belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if S can come to know on reflection that her belief that p is correctly inferentially based on her beliefs B<sub>1</sub> through B<sub>n</sub>.

This Access Condition constitutes a mixed or hybrid internalism since it requires the subject to have access to the fact that her belief is correctly inferentially based on its evidence but does not require her to have access to the fact *that* she has this access (she need not tell on reflection that she can tell on reflection that her belief is correctly inferentially based on its evidence). In consequence, it avoids Chisholm's mistake of requiring a subject to be able to tell on reflection that her belief is inferentially justified. But while (AC) possesses these virtues, it possesses two further features which some theorists would regard as problematic. I will describe these in turn.

First. As we will discuss more fully in Chapter 4, some theorists would prefer a more limited form of internalism. Relying on an intuitive distinction between *what*

---

<sup>25</sup>One question which I will not be able to take up here is whether the factors specified in the additional conditions discussed in section three below must also be reflectively accessible. I suspect that they need not be.

justifies a given belief and *how* what justifies does so, these theorists suggest that a plausible internalism need only require access to the former. With regard to a particular inferentially justified belief, *what* does the justifying are presumably the evidentiary beliefs on which that belief is based. *How* they justify, of course, is by serving as the basis for the subject's belief. Following Robert Audi, we may call the limited internalism which requires access only to what justifies a belief (i.e., the particular evidentiary beliefs on which it is based) *first-order internalism*, and we may call the more robust form of internalism which requires access also to *how* the evidentiary beliefs confer justification *second-order internalism*.<sup>26</sup>

The main reason why some internalists prefer to adopt the more moderate first-order position is that they are highly skeptical of ordinary people's abilities to determine on reflection how their beliefs receive justification from their evidence. In fact, some theorists even suggest that the relations whereby a belief receives justification from its evidence (e.g., relations of causal sustenance and evidential support) could *never* be reflectively grasped in the way that second order-internalism requires.

Despite these worries concerning second-order internalism, I do not believe that first-order internalism represents a preferable alternative in explicating what it means to be in a secure position with respect to the truth. As I shall argue in the fourth chapter, it does not make sense to require subjects to have access to their evidentiary beliefs without also requiring them to have access to at least some aspects of how these beliefs confer justification. It is difficult to understand how an individual's capacity to access her evidentiary beliefs could contribute to her being in a secure position with respect to the

---

<sup>26</sup>See Audi (1993), ch 11. If we interpret the notion of "how what justifies does so" loosely enough, the distinction between first and second order internalism can be regarded as exhausting the field of internalist theories. (Every justificatory factor is either part of *what* justifies, or involved in *how* what justifies does so.) Thus, any theory which requires more than what first-order internalism requires will be, on this suggestion, a variety of second-order internalism.

truth, unless she also has some awareness of how those evidentiary beliefs confer justification on the belief they support. For this reason, I prefer the initial formulation of the Access Condition given in (AC) to a weaker condition which would require only that the subject have access to the evidentiary beliefs themselves, but not to the fact that her belief is correctly inferentially based on them.

Second. Some theorists will balk at the requirement that a subject come to *know* that her belief is correctly inferentially based on its evidence (or that any relevant justificatory factors obtain). Robert Audi has suggested that it might be more plausible to treat reflective access to justificatory factors as a matter, not of knowing, but of *justifiedly believing* that the factors in question are present. After all, he argues, if internalism is a view about the requirements of justified belief, it seems unreasonable to require that a subject *know* that certain conditions or factors of justification obtain; the subject should only be required to believe with justification that these factors obtain.<sup>27</sup>

I disagree with this suggested weaker conception of reflective access (taking it to involve justified belief rather than knowledge) for the following reason. Working for the time being with the Access Condition in its initial formulation,<sup>28</sup> Audi would have us restate it so as to hold that a subject has an inferentially justified that *p* only if she can form on reflection a justified belief that her belief that *p* is correctly inferentially based on certain other beliefs. Now, the Access Condition imposes a requirement on inferential justification in the *categorical* sense, that is, a requirement on inferential justification *simpliciter*, as opposed to inferential justification to this or that degree. And, as I have argued previously, it is most natural to treat a claim that some belief is inferentially justified *simpliciter* as stating that the belief in question is inferentially justified to the

---

<sup>27</sup>Audi (1993), p. 344.

<sup>28</sup>The argument that follows will apply *mutatis mutandis* to the final version of the Access Condition given in Chapter 4.

*degree* required for knowledge.<sup>29</sup> It seems therefore that if a belief's being inferentially justified simpliciter requires that the subject be able to form on reflection a justified belief that her belief is correctly inferentially based on certain other beliefs, this second belief, to the effect that her inferentially justified belief is correctly inferentially based on others, must (when formed) be itself justified to the degree required for knowledge. Thus, even if we follow Audi in holding that a subject need not *know* on reflection that her belief satisfies certain justificatory conditions, but only that she be able to form on reflection a justified belief that her belief satisfies the relevant conditions, we must at least hold that this second belief is justified to the degree required for knowledge.

More than this is required, however, for the subject's "second-level" belief, formed on reflection, to the effect that her "first-level" belief is correctly inferentially based on others, must not only be inferentially justified to the degree required for knowledge, but must also be *true*. The second-level belief must be true because inferential justification requires that the subject's first-level belief actually *be* correctly inferentially based on others, and not simply believed with justification to be so based. Thus, having reflective access to justificatory factors or conditions involves, in the case of inferential justification, being able to form on reflection a *true* belief which is justified to the degree required for knowledge. In my view, if this much is required to have reflective access to justificatory factors or conditions, then it is altogether simpler to require a subject to *know* on reflection that the relevant factors or conditions obtain.<sup>30</sup>

---

<sup>29</sup>But, as I have said, this should not be construed to mean that the belief in question satisfies an "anti-Gettier condition," i.e., a condition which beliefs fail to satisfy in Gettier-style counter-examples to the traditional analysis of knowledge.

<sup>30</sup>Of course, the only thing that prevents a justified true belief that some justificatory factor obtains from automatically being an item of knowledge is the possibility of some Gettier-style situation occurring. It may be, however, that the sorts of "deviant causal chains" on which Gettier-type counter examples depend are not possible in cases of reflective knowledge.

The initial formulation of the Access Condition given above raises a number of issues which I will try to resolve in Chapter 4. To indicate the course of the discussion, I will briefly identify them.

First, the Access Condition as presently formulated makes it a requirement that the subject *can* come to know on reflection that her belief is correctly inferentially based on its evidence, but it is not clear how this modality is to be interpreted. It might be interpreted very weakly as requiring only that there be *some* logically possible circumstances in which the subject is able to tell on reflection that her belief is correctly inferentially based on others. On the other hand, it could be interpreted as a very strong requirement that the subject be able to tell this in *any* circumstance. Neither of these seems plausible, however. Thus, in Chapter 4, I shall argue for a middle course which requires that, where circumstances are *appropriate* for reflection, it is highly probable that the subject will come to know on reflection that her belief is correctly inferentially based on its evidence.

Second. Little has been said thus far about the meanings of terms such as “accessibility” and “knowledge by reflection.” The idea roughly is that a person has reflective access to the truth of a proposition *p* when she can come to know that *p* is true simply by focusing her attention on the question of whether that proposition is true. When characterized in this way, cases of reflective knowledge would appear to fall into two distinct categories. First, a subject would appear to have reflective access to simple truths about her present state of mind, e.g., that she is currently experiencing stabbing pain, or that she is currently imagining a purple beach ball, for she would appear to be able to come to know these things about herself simply by carefully focusing her attention on what she is currently experiencing or imagining. Second, a subject would also appear to have reflective access to simple truths of logic and mathematics, and other truths traditionally identified as knowable *a priori*, e.g., that  $2+2=4$ , or that nothing can be

red all over and green all over at the same time.<sup>31</sup> We might postulate two separate “faculties” of reflective knowledge, one an introspective faculty which yields knowledge of a subject’s own mental states, and the other a faculty of “rational insight” which yields knowledge of truths of reason and logical relationships.<sup>32</sup>

It would appear that both sorts of reflective knowledge are required if a subject is to have an inferentially justified belief. For knowing on reflection that a belief is correctly inferentially based on certain evidentiary beliefs would seem to require knowing on reflection both that it is evidentially supported by these other beliefs, and that it is causally sustained by them. A subject’s faculty of “rational insight” would presumably be responsible for the former sort of knowledge, and her faculty of introspection would be responsible for the latter.<sup>33</sup> One major question on which the viability of the Access Condition depends, then, is whether subjects are genuinely capable of having reflective knowledge of these relations.

Third. It might seem puzzling that one’s satisfying the Access Condition, as it is currently stated, *entails* that one satisfies the Basis Condition. This results from the fact that the former requires a subject to come to *know* on reflection that her belief is correctly inferentially based on others. But, of course, if she is able to know on reflection that her belief is correctly inferentially based on others then it must be *true* that it is correctly inferentially based on others, which means that the Basis Condition is satisfied. It will

---

<sup>31</sup>I do not mean to suggest that we should ignore the distinction between introspective and *a priori* knowledge. My point is merely that for our purposes it is useful to view these as two species of a more general variety of knowledge, namely, *reflective* knowledge.

<sup>32</sup>I borrow the term ‘rational insight’ from Bonjour’s (1985) and (1997) treatments of *a priori* knowledge, though I think he would regard the uses which I will eventually make of it as an abuse of this term.

<sup>33</sup>If evidential support relations are not logical, but rather nomological connections, it will have to be shown that knowledge by “rational insight” is not limited to truths of reason and logical relationships. I will try to indicate how this could be accomplished in Chapter 5.

emerge in the fourth chapter, however, that the initial formulation of the Access Condition is too strong. As I shall argue there, it cannot strictly speaking be a requirement of inferential justification that the subject come to know on reflection that her belief is correctly inferentially based on its evidence, since this is not something which can be *noninferentially* known, and (for reasons which will become clear) *only* justificatory factors which can be noninferentially known to obtain can be required to be reflectively accessible. In the end, then, what will be required to be reflectively accessible is not the fact that the subject's belief is correctly inferentially based on its evidence, but merely some of the factors which are *involved* in the belief's being correctly inferentially based on that evidence (specifically, the relations of evidential support and causal sustenance between the belief and its evidentiary beliefs). So, the Access Condition in its final formulation will not entail that the Basis Condition is satisfied.

At the end of the preceding section of the present chapter, I worried that some theorists would find the Basis Condition too implausible to be accepted. The same no doubt must be said for the Access Condition. In the first place, animals, small children, and conceptually unsophisticated adults will be unable even to raise the question of whether their beliefs are based on adequate evidence, let alone determine this on reflection. If the Access Condition is correct, none of these individuals will have inferentially justified beliefs. And given that a great many beliefs about the past, future, and things unobserved must be inferentially justified if justified at all, it follows that animals, small children and conceptually unsophisticated adults will have very few justified beliefs. This result alone would lead some theorists to reject the Access Condition.<sup>34</sup> But many would argue in addition that even adults who are fully capable of appreciating the question of whether their beliefs are based on adequate evidence will in general be incapable of answering this question on reflection. They claim that it is just

---

<sup>34</sup>See, e.g., Alston (1989), p. 240.

too much to expect that ordinary people be able to determine by reflection alone whether their beliefs are correctly inferentially based on evidence.

In Chapter 5, I will try to respond to these worries about the Basis and Access Conditions. First, I shall argue negatively that the degree to which ordinary people fail to satisfy my two conditions is exaggerated, and that in many cases people both hold their beliefs on the basis of adequate evidence, and are able to access that evidence on reflection. Second, I shall argue positively that, to the extent that the Basis and Access Conditions explicate the notion of a secure position, they offer a slightly idealized standard of justification which is useful for epistemic evaluation in those situations in which individuals *strive* to achieve this goal. These conditions need not be satisfied in every context, or by every creature, to be philosophically illuminating.

### 3. *The Question of Further Conditions.*

Putting together the two conditions so far discussed, we have the following partial account of inferential justification:

- (IJ3) S's belief that p is inferentially justified by her beliefs  $B_1$  through  $B_n$  only if
- (i) (*The Basis Condition*) Her belief that p is correctly inferentially based on her beliefs  $B_1$  through  $B_n$ , and
  - (ii) (*The Access Condition*) She can come to know on reflection that her belief that p is correctly inferentially based on her beliefs  $B_1$  through  $B_n$ .

In the next two chapters, I will develop this partial account in some detail. Before turning to this task, however, I should briefly touch on some further conditions which would need to be added to (IJ3) to obtain a complete theory of inferential justification. My aim is not to offer specific proposals in these areas, but to identify avenues of future research, and to point out certain issues which, though closely related to the topics of this dissertation, I will not have time to discuss more fully.

As I indicated in my initial discussion of (IJ1), the first condition which would need to be added concerns the epistemic status of S's beliefs  $B_1$  through  $B_n$ . If S's belief

that  $p$  is inferentially justified by her beliefs  $B_1$  through  $B_n$ , then this justification is, in effect, transferred from the latter beliefs to the former belief. But if a *transfer* of justification is occurring here, then the beliefs  $B_1$  through  $B_n$  must themselves be justified. If even one of these latter beliefs is not justified, then  $S$ 's belief that  $p$  is not genuinely justified. It is of no use to have beliefs which evidentially support the belief that  $p$ , and from which one has inferred that belief, if one does not in turn have evidence for the beliefs from which one did the inferring.

Therefore, (IJ3) needs to be supplemented with a "justification condition", that is, a condition requiring that the beliefs  $B_1$  through  $B_n$  be themselves justified.

Acknowledging this, however, does not get us very far, for the precise form which such a condition should take is in fact a very complicated and controversial matter. The principle question concerns whether the beliefs  $B_1$  through  $B_n$  must be justified *independently* of the belief that  $p$ .

On one view, if  $S$ 's belief that  $p$  is justified by her beliefs  $B_1$  through  $B_n$ , then the latter beliefs must be justified independently of the belief that  $p$ , meaning that they must in turn be supported by evidence which does not presuppose or include the belief that  $p$ . If one (or all) of the beliefs  $B_1$  through  $B_n$  includes in its evidence the belief that  $p$ , then, according to this view,  $S$  could not without begging the question base her belief that  $p$  on these other beliefs.

This view gives rise to a position traditionally known as *foundationalism*. According to this position, every inferentially justified belief ultimately depends for its justification on beliefs which are not inferentially justified. For, a given inferentially justified belief will depend for its justification on certain other beliefs, which on this view must be justified independently of the first belief. Their justification in turn may depend on a third set of beliefs being independently justified, and so on, thus generating a regress of ever more prior beliefs, that is, beliefs which must themselves be justified before that

justification can be transferred to succeeding beliefs. The claim is that, eventually, if one traces one's way back along the inferential chains which precede and support a given (justified) belief, one must in each case arrive at a terminating belief which is justified, but not in a way which depends upon further beliefs being justified. These noninferentially justified beliefs thus constitute the foundation on which the justification of each inferentially justified belief depends.<sup>35</sup>

An alternative view allows that a belief can be justified by other beliefs which are not, strictly speaking, justified independently of it. This is the position usually known as *coherentism*. On this view, one's justified beliefs form a "closed curve" so that, by tracing an inferential chain leading out from a given belief, one will eventually come back to the belief from which one started, giving the impression that the belief presupposes itself as a premise, or depends on its own prior justification. Foundationalists would say that the belief, in such a case, would not really be justified, because the subject is "begging the question" in holding it. The coherentist takes a different position, however. According to the coherentist, the foundationalist makes the mistake of thinking of justification in "linear" terms. The error, says the coherentist, lies in thinking that justification can be passed to a subsequent element in an inferential chain only if it has *first* or *already* been passed to the preceding element in that chain. Instead, the coherentist takes a nonlinear or "holistic" view of justification--one's entire web or system of beliefs is justified as unit, in virtue of logical, evidential, explanatory, probabilistic, in short, *coherence* relations among its member beliefs. The greater the number and degree of such coherence relations, the more one's system of beliefs is justified, and hence, the more one's individual beliefs are justified. On such a view *none*

---

<sup>35</sup>Contemporary foundationalists include Audi (1988) and (1993), Chisholm (1980) and (1989), Fumerton (1985), and Moser (1989).

of one's beliefs are truly justified independently of the others. Each belief depends on all of one's other beliefs for justification, and hence ultimately upon itself.<sup>36</sup>

So, there would appear to be two possible ways to supplement (IJ3) with a "justification condition." One possibility is to require that S's beliefs  $B_1$  through  $B_n$  be justified independently of her belief that  $p$ , thus yielding a foundationalist position. The other possibility is to allow the beliefs  $B_1$  through  $B_n$  to be justified in a way which does depend ultimately on the belief that  $p$ , thus yielding a coherentist position. Which of these formulations of the justification condition should be adopted is a question I will leave open in subsequent chapters. The result, I hope, will be an account of inferential justification which is compatible with both foundationalist and coherentist theories.

The second area in which further conditions of inferential justification still need to be worked out concerns the potential role of defeating counter-evidence in determining justification. Briefly, there are two sorts of cases that need to be considered.

First, there are cases in which a subject's belief that  $p$  is correctly inferentially based on evidence, but she has failed to recall additional evidence in her possession which "undercuts" that belief. The example discussed briefly earlier of S's belief that Jones would classify himself as politically conservative fits this description, since it is correctly inferentially based on its evidence (the evidentiary beliefs that Jones is a member of the National Rifle Association, and that ninety percent of the members of this association would classify themselves as politically conservative), but S has failed to recall the fact that Jones has voted for every Democratic candidate for President since George McGovern. Intuitively, the justification for S's belief is defeated by this counter-evidence. But does this mean that *any* counter-evidence stored in memory is relevant to the justification of a belief, or are there certain items of contrary evidence which, because

---

<sup>36</sup>Coherentist accounts of justification can be found in BonJour (1985), and Lehrer (1974) and (1990).

they are so deeply buried in long term memory, should not be thought to impugn the epistemic status of a belief if the subject fails to recall them? It is a question I must leave open.

Second, there is the sort of case, so far not discussed, in which a subject correctly infers a belief that *p* from the evidence in her possession, but where she *could have* gathered further evidence and, had she done so, she could not have correctly inferred that *p*. Imagine, for example, a manufacturer of rock climbing equipment who subjects a prototype harness to the standard battery of tests and concludes that the harness is safe. Suppose, however, that the manufacturer is aware that a new type of test has recently been developed, but because it is not mandatory to subject her product to this test, chooses not to conduct it. Imagine that, if she had subjected her prototype to this test it would have failed it in such a way as to make it strikingly clear that the harness is not safe. The issue is whether this potential evidence, which the manufacturer could have gathered, but chose not to, affects the epistemic status of her belief that her prototype harness is safe.

Like the debate between foundationalists and coherentists, the issues raised by the foregoing cases will have to be set aside. This means, admittedly, that important aspects of being in a secure position will remain unexplored in what follows (for the existence of potential counter evidence is clearly relevant to the security of one's position). But I think that the conditions which I will develop, i.e., the Basis and Access Conditions, provide the heart of what it is to be in a secure position with respect to the truth, and so the investigations of this dissertation may still throw considerable light on this important notion.<sup>37</sup>

---

<sup>37</sup>For an informative discussion of the role of defeating evidence in determining justification, see Feldman (1988a), (1988b), and (1992), Goldman (1993), pp. 9- 13, and Komblith (1983).

#### 4. Summary.

In this chapter, I sketched the following partial account of inferential justification, which I will develop over the next two chapters:

- (IJ3) S's belief that p is inferentially justified by her beliefs  $B_1$  through  $B_n$  only if
- (i) (*The Basis Condition*) Her belief that p is correctly inferentially based on her beliefs  $B_1$  through  $B_n$ , and
  - (ii) (*The Access Condition*) She can come to know on reflection that her belief that p is correctly inferentially based on her beliefs  $B_1$  through  $B_n$ .

In the third chapter I will focus on the Basis Condition component of (IJ3), and try to determine what factors beyond relations of evidential support and causal sustenance between a belief and its evidence are required in order for the former to be correctly inferentially based on the latter. In the fourth chapter I will examine the Access Condition component of (IJ3), and consider more carefully the question of which factors of inferential justification should be required to be reflectively accessible.

I also discussed in this chapter the principal difficulties which many theorists will have with the two conditions of (IJ3). It will be argued by these theorists that most ordinary people's beliefs are not causally sustained by evidentiary beliefs once they have been inferred, and that most ordinary people are not able to tell on reflection whether their beliefs are correctly inferentially based on others. As I have indicated, I will try to resolve these difficulties in Chapter 5.

## Three

### The Basis Condition

In the previous chapter, I discussed two conditions of inferential justification as a way of partially explicating what it is to be in a secure position with respect to the truth. The first of these, the Basis Condition, is the subject of this chapter. I examine the Access Condition in the next chapter.

Thus far, I have given only the following initial formulation of the Basis Condition:

(BC) S's belief that  $p$  is inferentially justified by her beliefs  $B_1$  through  $B_n$  only if her belief that  $p$  is correctly inferentially based on her beliefs  $B_1$  through  $B_n$ .

This initial formulation raises two questions with which the present chapter will be principally occupied. First, what is required, beyond a mere causal connection, for a belief to be *inferentially* based on other beliefs? Second, what is required for a belief to be *correctly* inferentially based on others? After resolving a few preliminary issues, I will do my best to answer these questions.

#### *1. Some Preliminaries.*

In this chapter (and in the dissertation generally), I will treat inference or reasoning as a process of belief expansion, that is, as a process which *generates* beliefs. The term 'inference' can also be taken to cover certain processes which lead to the revision or rejection of belief, but my focus will not be on inferences of these sorts. I will, in addition, treat inference as a *local* process, that is, as a process which operates, in the typical case, on a relatively small *subset* of a subject's total belief system. When, for example, an individual infers that the next swan she sees will be white, it seems most reasonable to regard her reasoning process as operating on, or taking as its "input," a

very small subset of her belief system, for instance, the subset consisting of her belief, say, that all swans are white. There may be instances of reasoning of a more global sort (as in the case, perhaps, of “paradigm shifts”?), but if there are, my account will, at least in its present form, ignore them.

The treatment of inference as a local process follows naturally from a certain view of the nature of belief which, as I shall be assuming it here, deserves some comment. I take it that when a subject believes or accepts a proposition, that proposition is in some sense stored or represented in her long term memory: the proposition is, as Gilbert Harman would say, “written down in mentalese.”<sup>1</sup> On this view, believing a proposition is not to be identified with having a disposition to *affirm* that proposition. For, the typical subject is disposed to affirm an infinite number of propositions, but, in view of the limited storage capacity of long term memory, has only a finite number of beliefs. Thus, I may have been *disposed* to affirm the proposition that there are fewer than 756 planets in our solar system, but prior to having explicitly considered that proposition (which I now have done), it was not represented or stored as a discriminable item in my long term memory, and thus was not a proposition I *believed*.<sup>2</sup>

---

<sup>1</sup>Harman (1986) p. 14. I do not know if Harman intends his reference to the “language of thought” to be taken literally. In any case, for our purposes, I believe, it may be taken figuratively. I do not mean to endorse, in other words, a computational theory of the mind.

<sup>2</sup>At any rate, it was not a proposition I believed in the sense in which I am interested. Harman (1986) distinguishes between *implicit* and *explicit* belief. The latter notion, that of a proposition’s being explicitly (and “assentingly”) stored in long term memory, is what I intend here in using the term ‘belief’. One merely implicitly believes a proposition when one does not explicitly believe it, but one would readily assent to it because it is easily inferable from other propositions that one explicitly believes (or because one is in some other way “committed” to it). The distinction which Harman draws between explicit and implicit belief corresponds roughly to Audi’s (1982) distinction between having a dispositional belief, and having a disposition to believe some proposition. (But see the discussion in footnote 4.)

At any given time, most of the beliefs stored in an individual's long term memory lie in what Alvin Goldman calls an "unactivated" or causally dormant state.<sup>3</sup> When a belief is called upon to play a role, either in guiding behavior, or in guiding via inferences the formation of new beliefs, it moves from an unactivated to an *activated* state.<sup>4</sup> In the typical case, only a few of the subject's beliefs can be activated at any one time, thus rendering her inferential processes more or less local in their operations.<sup>5</sup>

There is a view, which Goldman calls "inferential holism", which holds that *all* inferences are global, that every inference takes as premises all of the subject's prior beliefs, and has as its conclusion the resulting totality of beliefs. As a description of human inferential psychology, this view is not very plausible. As Goldman suggests, a subject can only reason from her *activated* beliefs, and it is very unlikely that a typical

---

<sup>3</sup>We should be careful to note, however, that unactivated explicit beliefs, while causally "dormant" in one sense, are nevertheless capable of *sustaining* other beliefs.

<sup>4</sup>Goldman (1986), pp. 200-201. Philosophers sometimes distinguish between dispositional and occurrent belief, where the former term implies belief stored in long term memory, while the latter implies belief that is currently entertained in consciousness, or presently "before the mind," as it were. Clearly, unactivated explicit beliefs are dispositional beliefs, but the relationship between activated belief and occurrent belief is less straightforward. A belief is active at any moment in which it is causally efficacious in guiding the subject's behavior or inferences. It seems to me that a belief could be causally active in this way without being *occurrent* in the sense just defined, so that a belief's being occurrent is not a necessary condition of its being causally active. (A subject can reason from premises some or all of which she does not consciously entertain.) For that matter, it is unclear to me whether a belief's being occurrent is *sufficient* for its being causally active. Accordingly, I will leave open the question of the relationship between activated and occurrent belief.

<sup>5</sup>As Goldman acknowledges, it is likely that belief activation is not an all-or-nothing matter, but admits of degrees, so that a given belief can be more or less activated than other beliefs at a particular time, and can smoothly fluctuate in its level of activation over time. For our purposes, it will be convenient to abstract away from this complexity, and treat belief activation as a simple on-off, or all-or-nothing matter. At the very least, it would seem that there must be some threshold of activation which a belief must attain before it can be causally efficacious in guiding behavior and inference. We might then treat beliefs which are at or above this threshold as activated *simpliciter*.

adult human being should be able, under any normal circumstances, to activate all of her beliefs simultaneously.<sup>6</sup>

In view of these facts, that inferences are in the typical case local, and that they operate only on activated beliefs, the issue of whether a given belief is or is not correctly inferred depends on two factors, namely, *which* of the subject's beliefs are causally active in the inference, and the precise character of the reasoning process which leads from those beliefs to the belief which is based on them. This means that the issue of defeating evidence discussed briefly in the last chapter, while relevant to the question whether a given belief is inferentially *justified*, is *not* relevant to the question whether that belief is correctly inferred *from* the beliefs which are active in producing it. This will greatly simplify the discussion that follows.

There is one final preliminary point that needs to be made. Thus far, I have spoken of inference as a process which generates belief. But I said in the previous chapter that inferential basing requires that a belief be *sustained* by the beliefs it is based on. I will proceed, therefore, on the following idealizing assumptions: First, that whenever a normal adult subject infers a belief from certain other beliefs, the inferred belief will subsequently be sustained by the beliefs it was inferred from. And, second, that a subject never holds a belief on the inferential basis of others unless she first inferred the former from the latter. Thus, ideally, a subject holds a belief on the inferential basis of other beliefs just in case it was inferred from those other beliefs, and is now causally sustained by them. While the actual story of inferential basing may be more complicated than this (given the fact that a belief, in reality, may be inferentially based on beliefs from which it was not originally inferred) it will for purposes of the present discussion be convenient to work within this idealized framework. At the end of the chapter I will say a

---

<sup>6</sup>Goldman (1986), pp. 206-207.

few words about how my account applies in cases where these idealizing assumptions are not satisfied.

In order to gain a fuller understanding of the nature of inferential basing, we must first focus more narrowly on the nature of *inference*, and the question of what distinguishes inferences from other sorts of belief generating processes. I will examine first the natural proposal, which I call the “rule-based” account, which holds that a belief generating process counts as an inference just in case it is governed by a *rule of reasoning* (or an “epistemic norm”, to use John Pollock’s term). I shall briefly criticize this proposal and then offer an alternative, based on suggestions from Barbara Winters and D. M. Armstrong. The alternative proposal is that reasoning is a process governed not by rules, *per se*, but by inferential *dispositions* or *habits* to form certain beliefs in the presence of other beliefs.<sup>7</sup> Of course, not just any belief forming habit or disposition will count as inferential, but only ones which meet certain further constraints. Winters and Armstrong claim that inferential dispositions must satisfy certain structural constraints which I will discuss below. My own view is that inferences are subject to a very different constraint. Following Robert Audi, I hold that inferences are habituated belief-state transitions which occur under the guidance of *connecting beliefs*. This means, roughly, that inferences are beliefs-state transitions which occur because the subject *takes* the belief she infers to be evidentially supported by the beliefs from which she reasons. Her taking the belief to be so supported *explains* why she comes to hold it. This is not to say that her reasoning must be fully articulate or conscious; the connecting belief requirement does *not* imply that inferences must in all cases be explicitly rehearsed. In fact, on Audi’s view, as I will discuss, a subject’s having the relevant connecting beliefs

---

<sup>7</sup>As I shall discuss below, it may be that an account of inference in terms of dispositions is translatable into an account in terms of rules, in which case the disposition-based account represents, not an alternative to the rule-based account, but a shift in emphasis from rules to dispositions as the most salient explanatory factor in inferential transitions.

need not even imply that, in any given case, she fully conceptualizes the evidential support relation in question, or even indeed *possesses* a concept of evidential support.

If inferences are habituated belief-state transitions which occur under the guidance of connecting beliefs, what then separates correct inferences from incorrect inferences? On the rule-based approach, an inference would be correct if it were governed by correct rules, and it would be incorrect if it were governed by incorrect rules (or by correct rules in an incorrect way). On the disposition-based approach, as it may be called, a habituated inferential transition is correct if it “corresponds to” or “tracks” a relation of evidential support, in the sense that it always leads from certain beliefs to a further belief which is evidentially supported by them.<sup>8</sup>

Given the question of whether inferences are governed by connecting beliefs, there are in fact two different approaches which an advocate of the disposition-based account of reasoning can take with regard to the nature of correct inference. On the one hand, there is what I call the “blind tracking” approach which requires only that a subject’s (correct) inferential processes track relations of evidential support without her necessarily having any awareness of the support relations those inferences are tracking. On the other hand, there is the “mediated tracking” approach, which I shall adopt, and which holds that the tracking of evidential support relations is, as the name implies, *mediated* by beliefs about those relations.

I turn now to examine and criticize the rule-based account of inference.

---

<sup>8</sup>I borrow the term ‘tracking’ from Nozick (1981). He uses it in his account of *knowledge*, whereas I employ it in my account of correct inference. While there is some similarity in our usage of this term, my understanding of it is, in the final analysis, very different from Nozick’s.

## 2. *Is Reasoning Governed by Rules?*

In this section I will consider and reject the view that reasoning is a rule or norm governed activity. I will look at two different interpretations of this so-called rule-based account, both of which, I believe, are unacceptable. The first interpretation holds that reasoning is rule governed in the sense that subjects deliberately apply or follow rules in reasoning. The second holds that rules govern reasoning only in the sense that subjects' reasoning processes automatically and habitually conform to patterns which fit, or can be described by, inferential rules.<sup>9</sup>

On the first way of treating the idea that reasoning is a rule governed activity, all human reasoning is taken to be analogous to the sort of activity a logic student undertakes in constructing, say, a formal proof in the sentential calculus. When a student constructs such a proof, she begins with premises which are "given" her, and proceeds to derive, or to try to derive, some assigned conclusion. This procedure involves several steps. From one or more of the given premises, the student derives a preliminary or intermediate conclusion. This intermediate conclusion then serves as a new premise from which (perhaps in conjunction with other intermediate conclusions, or in conjunction with some of the original premises) new intermediate conclusions may be drawn, and in this manner the desired final conclusion is eventually reached.

Each of these steps or transitions to new intermediate conclusions is licensed by some "rule of implication", the general form of such a rule being that if one has premises of a certain description D, then one is permitted to write down or accept on a new line a conclusion of description D'. These rules are *consciously* applied by the student at each

---

<sup>9</sup>One might reasonably ask why a third view is not mentioned, namely, that rules are causally operative in belief-state transitions (not merely descriptive of such transitions), without being consciously applied. I shall argue below that such an appeal to rules merely obscures the primary explanatory factor in inference, to wit, a subject's disposition to make belief-state transitions of a particular sort. On the view in question, inferential rules turn out to be merely descriptions of inferential dispositions, and as such are dispensable in favor of more direct talk about the dispositions themselves.

stage of the proof. The student recognizes that certain premises have the description appropriate for the use of some rule, and she draws the permitted conclusion. At each stage we can *explain* her behavior, her writing down a new line, in terms of her *realizing* that the line in question is licensed by a rule of implication in virtue of lines she has written down earlier in the proof.

Consider now a view which holds that ordinary human inference is analogous to the processes the logic student carries out in constructing her proof. On this view, if Jones sees (and thus believes) that the can of lighter fluid is very close to the flaming barbecue, and thus infers that he is in a dangerous situation, he reaches this belief by applying some rule that licenses certain belief-state transitions. Perhaps the rule in this case is the following: If one believes that one is in a situation in which a can of lighter fluid is very close to an open flame, then one is permitted to believe that one is in a dangerous situation. Such a rule, which tells one what beliefs are permitted in various circumstances, may be called, following John Pollock, an "epistemic norm." The suggested view, then, is that when Jones forms the belief that a can of lighter fluid is very close to an open flame, he somehow notices or apprehends that the epistemic norm in question applies in this case. Accordingly, he proceeds to deliberately apply this norm, and comes to believe that he is in a dangerous situation. Pollock calls this view, that reasoning involves the conscious and deliberate application of epistemic norms, the "intellectualist model of reasoning."<sup>10</sup>

This intellectualist model is implausible on a number of counts. First of all, if we suppose Jones to be an ordinary "person in the street," then it is unlikely that he operates with the *concept* of an epistemic norm, and thus could not consciously *apply* such norms in coming to hold beliefs.<sup>11</sup> Presumably Jones believes the *conditional*, that if one is in a

---

<sup>10</sup>Pollock (1986), p. 127.

<sup>11</sup>I take it that the *conscious and deliberate* application of an epistemic norm is a discriminative act, in that it implies an ability to recognize which cases are, and which are

situation in which a can of lighter fluid is very near an open flame, then one is in a dangerous situation. But this conditional is not itself an epistemic norm, for it does not specify conditions of permissible belief formation; indeed it makes no mention of the notions of belief or permissibility, nor in any other way makes a recommendation or prescription concerning reasoning or belief. Of course, if one *asked* Jones whether it is the case that if one believes that one is in a situation where a can of lighter fluid is very near an open flame, then one is permitted to believe that one is in a dangerous situation, he would no doubt acknowledge that such a principle seems eminently reasonable. But even this does not demonstrate that Jones has prior to this point operated with a worked out concept of epistemic norms which would enable him to apply such norms in forming beliefs.<sup>12</sup>

But, second, even if Jones *does* have the concept of an epistemic norm, and does at times deliberately apply norms in drawing conclusions, certainly the situation imagined is not one in which it is *necessary* for Jones to deliberately apply a norm in coming to believe that he is in a dangerous situation. He will presumably just come to have this belief, without wasting time applying any norms. Thus, at the very least, the intellectualist model is not an accurate account of *all* instances of human reasoning, and most likely is not even accurate with respect to the majority of instances.

---

not, subsumed under that particular norm. It seems to me, therefore, that one cannot consciously and deliberately apply epistemic norms in the manner required by the intellectualist model of reasoning unless one operates with the concept of an epistemic norm. This is not of course to say that the application of epistemic norms requires that one be familiar with the *term* 'epistemic norm.' Rather, such application requires that one operate with the idea that certain beliefs are *permissible* if one holds certain other beliefs.

<sup>12</sup>More exactly, I should say that it does not show that he has prior to this point operated with a concept of an epistemic norm which would allow him to apply such norms in the manner required by the intellectualist model of reasoning, that is, consciously and deliberately. I have thus far spoken as if this is the only manner in which epistemic norms might be applied. But shortly I will take up the suggestion that inferential transitions might be governed or guided by norms *without* being consciously and deliberately applied.

Finally, as Pollock points out, the intellectualist model seems to lead to an infinite regress. For the deliberate application of an epistemic norm would itself appear to be a kind of *reasoning* in that the subject subsumes a particular case, namely, a concrete situation in which she holds certain beliefs, under a general rule to the effect that in such situations it is permissible to draw such and such a conclusion. But if the intellectualist model is correct, the application of this epistemic norm, inasmuch as it constitutes an episode of reasoning, must in turn be governed by a *second* norm, a *meta-norm*, as it were, to the effect that the first norm may be applied in the situation in question. Thus, if Jones applies the norm, call it N, that if one believes that one is in a situation in which a can of lighter fluid is near an open flame, then one is permitted to believe that one is in a dangerous situation, then it seems that his reasoning in applying this norm is governed by the following meta-norm: If one believes that one is in a situation where a can of lighter fluid is near an open flame, then one is permitted to apply norm N. But if this meta-norm is in turn consciously applied, as it must be on the intellectualist model (since epistemic norms govern reasoning, on this view, by being consciously applied), there needs to be yet a *third* norm which governs the second norm's application, and so on. This regress cannot continue indefinitely--it is clear that if reasoning is governed by epistemic norms, then at some point they must govern reasoning without being consciously applied by the subject. Thus the intellectualist model is false as a general account of human reasoning.<sup>13</sup>

Pollock suggests an alternative to the intellectualist model of reasoning. On his view, reasoning is typically automatic and habituated behavior, not deliberate and consciously rehearsed, but behavior which is nevertheless governed by norms. He assimilates reasoning to other sorts of automatic and habitual behavior such as that exhibited by someone who has mastered bike riding or tennis playing. According to

---

<sup>13</sup>Pollock (1986), pp. 127-128.

Pollock, bike riding and tennis playing are in some sense norm governed activities. Certain rules apply which tell what one may do, or should do, in a given situation. Thus, there is the norm governing bike riding which says that if one feels the bicycle leaning to the right, then one should turn the handle bars to the right. And when playing tennis, there are, for example, rules which dictate how one ought to hold one's racquet when attempting a backhanded return. These rules may be consciously followed or applied by the novice, but once they are thoroughly mastered, they govern bike riding or tennis playing behavior without individuals having to think about them at all. How is it possible for a norm to govern behavior, on Pollock's account, if it is not consciously applied or followed?

"What makes it possible," says Pollock, "for your bike riding behavior to be governed by norms without your thinking about them is that you *know how* to ride a bicycle."<sup>14</sup> On Pollock's view, one *knows how* to perform an activity when one has internalized the rules or guidelines governing that activity, when one's performance of the activity naturally and automatically conforms to those rules without one's having to think about them. According to Pollock, something similar applies to reasoning, except of course that, in the case of reasoning, many of the norms or guidelines are prewired, part of a subject's innate cognitive endowment, and may never have entered her sphere of explicit awareness. The essential principle is nevertheless the same:

We know how to reason. This means that under various circumstances we know what to do in reasoning.... Our epistemic norms are just the norms that describe this procedural knowledge.... They describe an internalized pattern of behavior that we automatically follow in reasoning, in the same way that we automatically follow a pattern in bicycle riding.<sup>15</sup>

Focus on Pollock's saying that epistemic norms "describe a pattern of internalized behavior." The suggestion is that our reasoning is rule governed, because our reasoning

---

<sup>14</sup>Ibid., p. 129.

<sup>15</sup>Ibid., p. 131.

conforms to a pattern which is described by a set of rules. We have certain habits of belief formation, and these habits are governed by rules *in the sense that they are described by rules*.

The problem with this suggestion is that if rules merely *describe* inferential transitions, then they play no *explanatory* role in the psychology of inference. If Pollock is right, then Jones does not infer that he is in a dangerous situation *because* of a rule which permits this belief, but because he is disposed automatically to come to believe that he is in a dangerous situation when he believes that he is in a situation in which a can of lighter fluid is very near an open flame. The rule enters into the inferential transition only as a way of *describing* the transition: we say that Jones is disposed to reason in a way that conforms to the rule that if one believes that one is in a situation in which a can of lighter fluid is very near an open flame, then one is permitted to believe that one is in a dangerous situation. But all this really means is that when Jones has the belief specified in the antecedent of the rule, he is disposed to form the belief specified in the consequent of the rule. In that case, why bother talking about the rule? Why not focus one's attention instead on the *inferential disposition* as the truly salient factor in the inferential transition?

Perhaps Pollock could suggest that we need to appeal to rules in the account of *correct* inference, that correct inference is inference which conforms to a correct rule. Jones's inference, for example, is correct because the rule that if one believes that one is in a situation in which a can of lighter fluid is very near an open flame, then one is permitted to believe that one is in a dangerous situation, is (in some sense) a correct rule of inference. Thus, while the rule may not play any genuine role in explaining Jones's inference, it plays an epistemological role in explaining why Jones's inference is correct. But, as we shall see below, it is possible to construct an adequate theory of correct inference which appeals to properties of inferential dispositions, not to rules. On the view

which will be developed in section five, the correctness of Jones's inference primarily depends on two factors: First, his belief that he is in a situation in which a can of lighter fluid is very near an open flame evidentially supports the belief that he is in a dangerous situation.<sup>16</sup> Second, his inferential disposition *tracks* this relation of evidential support. That particular disposition will never yield a belief which is not evidentially supported by the belief (or beliefs) which, as it were, "trigger" the disposition and initiate the inferential process. We can therefore explain the distinction between correct and incorrect inference without appeal to the notion of a rule, but rather in terms of the notions of inferential disposition and evidential support.<sup>17</sup>

To summarize, the problem with Pollock's version of the rule-based account is that on his conception rules merely describe our inferential transitions, but play no explanatory role in the psychology of inference. What is explanatorily salient are just the dispositions which subjects have to make inferential transitions of certain characteristic sorts. But if inferential dispositions, and not inferential rules, do all of the work in explaining our inferences, then the rules can be effectively ignored in theorizing about the

---

<sup>16</sup>Perhaps the former belief evidentially supports the latter only when taken in conjunction with certain general principles or laws of nature governing the behavior of liquids such as lighter fluid in proximity to open flames. We might postulate that Jones believes the general principles, and that these beliefs are causally operative in his reasoning, although some might regard this as implausible, particularly if Jones is a relatively unsophisticated reasoning subject. Alternatively, we could take the view that there is a relation of evidential support between the belief that he (Jones) is in a situation in which a can of lighter fluid is very near an open flame and the belief that he is in a dangerous situation if the relevant general principles are *true* (and that there is no such relation if they are not). Such a view is closely related to the so-called "presupposition theory of induction". (See, for example, Burks 1963, and Nozick 1981, p. 263.)

<sup>17</sup>One could try to argue that the notion of evidential support must be understood in terms of rules of inference. But as we discussed in Chapter 1, evidential support may well be a factual, objective (i.e., either logical or nomological) connection, not to be identified with or explained in terms of rules. Consider again the analogy between evidential support and entailment. The premises that all humans are mortal and that Socrates is human entail the conclusion that Socrates is mortal. There are also various rules which license inferring the latter from the former. But the relation of entailment is not to be identified with, or even explained in terms of, these rules.

nature of inference. We can construct good theories both of the nature of inference, and of the nature of correct inference, without reference to inferential rules. The rule-based account of inference thus gives way to a disposition-based account of inference. One could, of course, try to develop a view in which inferential rules appear to play an explanatory role in inference. One could always claim, for example, that in Jones's case, he infers that he is in a dangerous situation *because* such an inference accords with the rule that if one believes that one is in a situation in which a can of lighter fluid is very near an open flame, then one is permitted to believe that one is in a dangerous situation, thereby apparently attributing to the rule an explanatory role. But what would this really mean except that Jones is disposed to reason in a way which can be described by the rule in question? Consider, as an analogy, a plant which is disposed to turn toward a light source when illuminated from a particular direction. One could try to explain the plant's behavior by saying that the plant turns toward the light because in doing so it satisfies the rule that if one is exposed to light from a certain direction, then one is permitted to turn toward the light. But this rather odd way of describing the plant's behavior merely obfuscates the principal explanatory factors in the situation, namely, that the plant is disposed to turn toward the light, and that the plant is indeed currently exposed to light. Similarly, an account of inference in terms of rules distracts attention away from what I believe is the primary explanatory factor in inferences, namely, the inferential disposition to form a belief when prompted by certain other beliefs.

I conclude that the rule-based account of inference gives way to a disposition-based account which, I believe, more accurately captures the relevant explanatory factors in human inference. The problem now is to develop such an account in detail.

### 3. *The Disposition-Based Account of Inference.*

Inference is a causal notion. If one infers a belief from certain other beliefs, then these other beliefs are the cause of the belief that one infers. But the mere existence of a causal connection between beliefs is not sufficient for inference. The causal connection between an inferred belief and the beliefs from which it was inferred must depend on, or derive from, a disposition on the subject's part to make inferential transitions of just that sort. An inference is a manifestation of an inferential disposition.<sup>18</sup>

From the fact that a glass is fragile, it does not follow that every possible way in which it might shatter would be a manifestation of that fragility. If the glass breaks because it is dropped on the floor, then this is a manifestation of its fragile disposition. If the glass breaks because it is struck by a large wrecking ball, then its breaking is not a manifestation of its fragility. Many things which are not fragile would break upon being struck by a wrecking ball.

By the same token, just because one is caused by certain of one's beliefs to form some additional belief, it does not follow that the latter resulted from *one's disposition* to make belief-state transitions of a particular sort. Thus, borrowing an example from Barbara Winters, I may believe that I am late for an appointment, and as a result become so nervous that my hands begin to shake, whereupon I may come to believe that I have palsy. So, my belief that I am late for an appointment causes me to believe that I have palsy, but I clearly do not *infer* the latter from the former, for I lack a *general* disposition to believe that I have palsy whenever I believe that I am late for an appointment (it is only in this particular instance that the one belief happens to yield the other).<sup>19</sup> Or, to take an example from Armstrong, my belief that someone is in the house might cause me to open the door, which in turn might cause me to believe that no one is in the house. But I

---

<sup>18</sup>The reader should not worry about the threat of circularity lurking in this statement. Inferential dispositions will not be *defined as* those dispositions which yield inferences.

<sup>19</sup>See Winters (1983).

am not generally disposed to believe that no one is in the house whenever I believe that someone is in the house. If I am, on a certain occasion, talking to someone in the house, it would be awfully strange if I were to form the belief that no one is in the house.

These examples illustrate the usefulness of the notion of inferential dispositions in understanding inference. It is in fact quite difficult to construct an adequate account of inference which successfully excludes the foregoing cases, but which makes no mention of inferential dispositions. For instance, in *A Materialist Theory of the Mind*, in which the second of the two preceding examples was first offered, Armstrong presented the following account of inference:

- (A1) S infers p from q just in case:
- (i) S's believing q *causes* him to acquire the belief that p, and
  - (ii) The belief that q is causally relevant right up to the time that S acquires belief p. (Unless S had believed q up to the time of acquiring belief p he would not have come to believe p).<sup>20</sup>

Condition (ii) was meant to exclude as inferences cases of the sort just discussed. Thus, once my belief that I am late for my appointment has caused my hands to shake, it is no longer causally relevant to my forming the belief that I have palsy. I could cease to have the former belief upon seeing my hands shake, and I would still come to hold the latter belief. A similar point applies to the second case mentioned above. Once I begin turning the doorknob, I could absentmindedly forget why I decided to check inside the house, and I would still come to believe that no one is in the house. So, my belief that someone is in the house is not causally relevant up to the time I form my belief that no one is in the house.

---

<sup>20</sup>See Armstrong (1968), 194–196. Armstrong does not offer his account in the tidy package presented here, but the statements of his two conditions are quoted directly from the text (except that I have used 'S', where he uses 'A'). See p. 194 for condition (i), and p. 196 for condition (ii). I should note also that (A1), strictly speaking, gives only Armstrong's 1968 account of "direct inference," the complete view being that an individual infers a belief just in case it is either directly inferred, or arrived at through a series of direct inferences. This complication will not effect the discussion, however.

Therefore, (A1) can explain why these two cases are not inferences. But (A1) is nevertheless flawed precisely because it fails to incorporate the notion of an (inferential) belief-forming disposition. Armstrong himself was at the time aware of counter-examples to (A1). For instance, he discussed the following case:

Suppose... that I believe p. Suppose, further, that this belief does not remain unconscious, but I introspectively acquire the belief that I believe p. Here the belief-that-p causes me to acquire the belief that I believe p. Now, nobody would want to call this a case of inferring that-I-believe-p from p,

but, unfortunately, it fulfills the conditions set down in (A1).<sup>21</sup> The belief that p causes, via my introspective faculties, the belief that I believe that p, and the former is presumably causally relevant up to the time that I arrive at the latter belief. In response, Armstrong could only say:

I do not know what to say about [this] case, and can only hope, like Hume on a celebrated occasion, that the instance is “so particular and singular” that it “does not merit that for it alone we should alter our general maxim.”<sup>22</sup>

But the fact of the matter is that the case is not “particular and singular.” We routinely come by introspection to hold “meta-beliefs” about what beliefs we do or do not hold. Since we ought to be able to distinguish between beliefs formed by introspection and beliefs formed by inference, (A1) cannot be an acceptable account of the latter.

To re-emphasize the problem, (A1) fails to explicate inference in terms of the notion of an inferential disposition. To understand inference, it is not enough to appeal to notions of causation and “causal relevance.” We have to think of the belief as resulting from a certain disposition on the part of the subject to form beliefs of that sort when given other appropriate beliefs as “inputs.”<sup>23</sup>

---

<sup>21</sup>Armstrong (1968), p. 200.

<sup>22</sup>Ibid.

<sup>23</sup>But, as I will discuss further below, the mere fact that a belief results from a subject’s disposition to make belief-state transitions of a particular sort is not *sufficient* to make it

What does it really mean to say that a given inferential transition results from a habit or disposition? It means, first of all, that the transition fits a repeatable pattern, and that the subject is so constituted that, under the right circumstances, the pattern will be repeated. Thus, if, upon checking every item in my shopping cart, I come to believe of each that its cost is under one dollar, then, provided that I believe that I have checked every item, I may come to believe that all of the items in my shopping cart cost under one dollar. This transition forms a pattern which I am disposed to repeat in the right circumstances. If I come to believe of each item in a given class that it has a certain property F (and I believe that I have observed every member of the class), I will in general be disposed to believe that all of the items in that class have the property F.

Armstrong himself later gave up the account given in (A1), and explicitly acknowledged the role of dispositions in inference. Thus, in *Belief, Truth, and Knowledge*, we find the following account, consisting of two conditions, which Armstrong calls the “dispositional condition” and the “manifestation condition”:

(A2) S infers her belief that p from her belief that q just in case:<sup>24</sup>

*The dispositional condition*

- (i) There exists some general proposition (x)(if Fx, then Gx), such that q has the form Fb, and p has the form Gb, and such that S is disposed so that: if S believes something of the form Fx, then this belief-state will both create (if necessary) and weakly causally sustain<sup>25</sup> within S's mind the belief that the corresponding proposition of the form Gx is

---

an inferred belief. Introspective beliefs, quite plausibly, result from *introspective* dispositions. Thus, inferred beliefs must result from inferential dispositions.

<sup>24</sup>In actuality, Armstrong speaks here of q being one of S's “putatively conclusive reasons for believing that p” (once again, I have replaced ‘A’ with ‘S’). I think it is safe to recast his view in terms of inference. By the phrase ‘a putatively conclusive reason’, Armstrong does not mean a (putatively) *logically* conclusive reason, but merely a reason which is putatively one which is sufficient to warrant, or justify, or make rational, the acceptance by the subject of the belief in question.

<sup>25</sup>According to Armstrong, one state of affairs *weakly* causally sustains another if it is by itself sufficient to maintain that state of affairs, but is not necessary, because the sustaining of the state of affairs in question may be overdetermined. Since, in the context of this discussion, I am concerned with inference as a belief *generating* process, I will, for my own part, ignore the question of what sustains a belief after it has been inferred.

true.

*The manifestation condition:*

- (ii) Baq both creates (if necessary) and weakly causally sustains Bap within S's mind.<sup>26</sup>

My stance toward (A2) is that it is *approximately* correct, but that it needs to be reformulated in certain respects. First of all, the manifestation condition, as Armstrong calls it, is not quite properly stated. For it does not follow from the fact that "Baq" creates "Bap" that it does so by way of any disposition on S's part. The manifestation condition should therefore explicitly state that S's belief that q "triggers" S's disposition to form the belief that p.

A second misgiving which I have about (A2) concerns the "general proposition" mentioned in the disposition condition. Armstrong claims that with "judicious torture" most inferential patterns can be made to fit the form ' $(x)(\text{if } Fx, \text{ then } Gx)$ .'<sup>27</sup> But it seems to me desirable to impose a less specific restriction on the form of inferential patterns. Thus, in place of Armstrong's general proposition, let us introduce the following quasi-technical notion: Let us say that when a subject draws an inference there is a *sequence* of beliefs, the first members of which are the "premises" of the inference, the final member being the "conclusion." In an inference, this sequence is of a general kind K such that the subject is disposed to form other belief sequences of the same kind. I propose then, to replace (A2) with the following preliminary account:

- (I) S's belief that p is inferred from her beliefs B<sub>1</sub> through B<sub>n</sub> just in case:
- (i) The sequence of beliefs  $\langle B_1, \dots, B_n, p \rangle$  is of a kind K such that S has the following disposition D: If S has the first n beliefs of a sequence of kind K, then she is disposed to form the final belief in that sequence, and
  - (ii) S's coming to believe that p is in fact a manifestation of disposition D, triggered by her beliefs B<sub>1</sub> through B<sub>n</sub>.

---

<sup>26</sup>Armstrong (1973), p. 85. By 'Bap,' Armstrong means the state of S's believing that p (similarly for q).

<sup>27</sup>Ibid., p. 86.

But even this proposal is imperfect. First, I have no specific guidelines as to how the “kind K” is to be determined in a given case. Hopefully, the relevant kind will be obvious in many circumstances. To return to the example of the items in my shopping cart, the relevant kind K in that case would seem to consist of sequences of beliefs the first members of which assert of each item in a certain class that it has a property F (and that every item in the class has been observed), and the final member of which asserts that all of the items in that class have the property F. So, in many instances, K will be specified in terms of the contents of the beliefs involved in the relevant sequences. In other instances it might be appropriate to speak of K being specified in terms of the *form* of the beliefs involved (e.g., the case where one infers a belief that q from beliefs that p, and that if p, then q).<sup>28</sup>

Second, there is an immediate problem with the way that condition (i) has been written. As it stands, the condition says that if the subject S has the first n beliefs in a sequence of kind K, then she is disposed to form the final belief in that sequence. But there may be several sequences of kind K which share the first n beliefs, and there may thus be no unique final belief which S could come to infer. For example, S might believe that the President likes football, and come thereby to believe that either the President likes football or the President likes baseball. In this case, K would consist of 2-membered sequences, the second member of which is a disjunction, one of whose disjuncts is the first belief in the sequence. There will then be an infinite number of sequences of kind K each having as its first member the belief that the President likes football. In response to this observation, condition (i) of (I) should be reformulated in the following way:

- (i') The sequence of beliefs  $\langle B_1, \dots, B_n, p \rangle$  is of a kind K such that S has the following disposition D: If S has the first n beliefs in a sequence of kind K, then she is disposed to form an additional belief which would complete a sequence of kind K.

---

<sup>28</sup>In what follows I will ignore the problem of determining the relevant kind K, and assume that in each case it is intuitively obvious.

Third. (I) is problematic also in that, in general, one will not be disposed to form the final belief in a sequence simply because one has the first beliefs in that sequence. For example, Goldman discusses the case of bacteriologists who, before the discovery of penicillin, knew that molds caused clear spots in bacteria cultures, and also that such clear spots indicated an absence of bacterial growth, but failed to draw the obvious conclusion that molds release an anti-bacterial agent. "They apparently failed to activate these beliefs at a single time, juxtapose them, and then make a natural inference."<sup>29</sup> This suggests that we must modify the first condition in (I) yet again:

- (i'') The sequence of beliefs  $\langle B_1, \dots, B_n, p \rangle$  is of a kind  $K$  such that  $S$  has the following disposition  $D$ : If  $S$  has the first  $n$  beliefs in a sequence of kind  $K$ , and these beliefs are simultaneously activated, then she will be disposed to form an additional belief which completes a sequence of kind  $K$ .<sup>30</sup>

But, even with this revamped condition (i'') plugged into the original proposal (I), the latter is still unacceptable, I believe. The problem is that many dispositions of the sort described in (I) do not qualify as *inferential* dispositions. Consider, for example, an individual who is deranged in such a way that whenever he believes that there are at least three people in the room, he comes also to believe that the President of the United States is in the room. There would seem to be a kind  $K$  of (in this case, 2-membered) belief sequences, such that whenever the individual in question comes to believe the first member of the sequence, he is disposed to come to believe the second member. And we may suppose that, on any given occasion, his coming to believe that the President of the United States is in the room is a manifestation of this disposition. But this individual

---

<sup>29</sup>Goldman (1986), p. 203- 204. I follow Goldman in distinguishing between activated and unactivated beliefs. An unactivated beliefs lies "dormant" in long term memory; it is activated when it becomes causally operative in guiding behavior or in generating inferences. Whether a belief must be *occurent*, or manifested in consciousness, in order to be thus active is a question I will not try to answer.

<sup>30</sup>Condition (ii) must also be modified to indicate that it is the *activation* of the beliefs  $B_1$  through  $B_n$  which triggers the formation of the belief that  $p$ .

does not *infer* the belief that the President of the United States is in the room from his belief that there are at least three people in the room.

One *could* try to claim that this is in fact a case of inference, albeit a *deranged* sort of inference. But why think of this particular belief-state transition, from the belief that there are three people in the room to the belief that the President of the United States is in the room, as an inference? It seems more reasonable to think of it just as a strange and sad condition to which the individual is subject, that he always comes to believe that the President is in the room whenever he believes that there are three people in the room.<sup>31</sup>

Or, consider once again the introspective example discussed above. Suppose I believe that the President prefers baseball to football, and, by reflectively introspecting the matter, I come to believe that I believe this. Here my belief that I believe that the President prefers baseball to football *does* result from a disposition, but not an *inferential* disposition, although, again, (I) seems to class this as a case of inference. For, there is a kind K of belief sequences such that, under the right conditions, I am disposed to hold the second member when I hold the first member of the sequence (specifically, when I believe that p, I am disposed, at least under the right conditions, to believe that I believe that p).

Finally, the example considered at the beginning of this section, in which my belief that I am late for an appointment causes me to believe that I have palsy, can be refashioned into a counter-example to (I). Focus not on the second belief, that I have palsy, but on the (presumably intermediary) belief that my hands are shaking. I may well be so disposed that when I believe that I am late for an appointment, my hands begin to shake, and so disposed that when I my hands begin to shake, I believe that they are

---

<sup>31</sup>However, as we shall see below, this could indeed count as a case of inference if a certain further condition is satisfied, namely, that the transition from the belief that there are three people in the room to the belief that the President of the United States is in the room occurs under the guidance of an appropriate "connecting belief."

shaking. My belief that my hands are shaking thus results from a general disposition of a content-specified kind K (in particular, the disposition to believe that my hands are shaking whenever I believe that I am late for an appointment), and this disposition is triggered by the activation of my belief that I am late for an appointment. The conditions of (I) are then satisfied, but intuitively I do not infer my belief that my hands are shaking from my belief that I am late for an appointment.<sup>32</sup>

So, (I) requires modification yet again, if these cases are to be excluded as inferences. One possible way of handling cases of these sorts is suggested by Barbara Winters:

(W) S infers q from other beliefs p only if the set of beliefs [p, q] instantiates an inference pattern and S's transitions from p to q is a result of S's general disposition to make transitions that exhibit that form.<sup>33</sup>

Of interest, for the present moment, is the claim that the set of beliefs [p, q] must "instantiate an inference pattern."<sup>34</sup> According to Winters, an "inference cannot take place unless the relevant beliefs are *connectable*, related in such a way that a belief may be concluded from them.... Reasoning cannot occur among 'unrelated' beliefs."<sup>35</sup> She says, in addition, that "we can only view a set of beliefs as an inference if we can understand it as instantiating some sort of inference pattern, rational or irrational...", and, finally, she claims that the reason why, for example, my belief that I have palsy (acquired

---

<sup>32</sup>Notice that quite plausibly I do infer my belief that I have palsy from my belief that my hands are shaking (together perhaps with a premise to the effect that the best explanation for my shaking hands is that I have palsy).

<sup>33</sup>Winters (1983). Once again, I have replaced 'A' with 'S'.

<sup>34</sup>Winters speaks of the *set* of beliefs [p, q]. The reason I would prefer instead to speak of the *sequence* of beliefs <p, q> is that the *order* of the beliefs matters. There is a belief which is being inferred and a belief (or beliefs) from which one is doing the inferring. The set of beliefs [p, q] is identical to the set [q, p]. But <p, q> and <q, p> are different sequences. I will not quibble over this point, however.

<sup>35</sup>*ibid.*, p. 212.

via my belief that I am late for my appointment) does not constitute a case of inference is that the beliefs involved “fail to instantiate a recognizable pattern.”<sup>36</sup>

Winters does not tell us what constitutes a “recognizable inference pattern,” but her general intention is clear enough. There is *no* such pattern in the case of the individual who comes to believe that the President is in the room whenever he believes that at least three people are in the room; there is no such pattern in the case of my belief that I believe that the President prefers baseball to football, arrived at introspectively from my belief that the President prefers baseball to football; and there is no such pattern when I come to believe that my hands are shaking as a result of my belief that I am late for an appointment.

It is important to realize that Winters’s “recognizable inference pattern” requirement does not imply that *only* reasonable or rational inferences are possible. Thus, if I come to believe that the cat knocked over the milk carton, as a result of my belief that the milk carton is spilled on the floor, then we might allow that this transition instantiates a “recognizable pattern” despite the fact that it may be an unreasonable inference (because, e.g., both the dog and the cat were in the kitchen when the carton spilled). To put this point another way, the recognizable inference pattern requirement is not meant to imply that an inference occurs only if the preceding beliefs in the transition genuinely evidentially support the subsequent belief. Perhaps we can say that, on Winters’s view, inference requires only that there be a relationship between the premises and the conclusion which, as it were, “mimics” a relation of evidential support. Or, even better, on Winters’s view the premises in an inference must *at least purport* to evidentially support the conclusion drawn from them.

This last point, I believe, captures the essence of inference. In an inference the premises purport at least to evidentially support the conclusion, or, instead, the subject in

---

<sup>36</sup>*Ibid.*, pp. 214-215.

drawing the inference purports to reach a conclusion which is evidentially supported by her premises. This is why the deranged individual who comes to believe that the President is in the room because he believes that three people are in the room does not infer the former belief. He does not purport to reach a belief which is evidentially supported by premises in his possession. He simply believes that the President is in the room because he believes that three people are in the room, and that is the end of the story.

While I believe that Winters' recognizable inference pattern requirement is on the right track, in that it suggests the view that inference involves purporting to reach a conclusion which is evidentially supported by one's premises, I do not believe that the account presented in (W) will ultimately work, for two reasons. First, I believe that there are examples of inference in which no recognizable pattern exists. I think that *under the right conditions*, even as strange a transition as believing, say, that the moon is made of green cheese, because one believes that Socrates is mortal, could count as an inference.<sup>37</sup> Second, to say that a pattern must be recognizable raises the obvious question "Recognizable to whom?" Must the pattern be *generally* recognizable? Or, recognizable to the majority of human beings? Or, the majority of adult human beings? Or, perhaps, the majority of adult human beings educated in such matters? What if the pattern is recognizable to the subject, but not to any one else? What if there are extra-terrestrial beings, creatures living on a planet in orbit around Vega, say, who reason according to inference patterns so different from ours that we are unable to recognize their patterns as inferences? Does it follow that such creatures do not draw inferences? Such questions (and the absence of answers) suggest to me that it is best not to impose on inference a recognizable pattern requirement.

---

<sup>37</sup>This example is borrowed from Armstrong (1968). As mentioned in footnote 31, seemingly bizarre transitions of this sort may nevertheless count as inferences if they occur under the guidance of appropriate connecting beliefs. See below for further details.

Still, when a subject draws an inference, she purports to arrive at a belief which is evidentially supported by her premise beliefs. How do we explain this if not in terms of a recognizable pattern which her belief-state transition instantiates? The answer is that the subject purports to draw an evidentially supported conclusion in that she *takes* her conclusion to be supported by her premises. An inference, that is, is a belief-state transition which occurs under the guidance of what Robert Audi calls a *connecting belief*.

#### 4. *The Role of Connecting Beliefs.*

In the previous section, I rejected Barbara Winters's view that inferential dispositions must fit recognizable inference patterns. It is true that an inference must be, in some sense, *rationaly intelligible*, but not because it fits a pattern which a third party would recognize as inferential. A subject's belief-state transition is intelligible as an inference if she is led to her conclusion because she *thinks* (rightly or wrongly) that it is evidentially supported by her premises. A belief-state transition, no matter how bizarre it might otherwise seem, surely *makes sense* (we understand why the subject engaged in that particular transition) if it occurs in part because the subject believes that her prior beliefs evidentially support the belief she subsequently forms. So, an inferential disposition need not fit a recognizable pattern, but it must be triggered by a connecting belief.

This suggestion, that inferential dispositions must be triggered by connecting beliefs, gains support from some of the cases discussed in the previous section. Thus, with regard to my forming the belief that my hands are shaking as a result of my belief that I am late for my appointment, it might plausibly be suggested that one salient factor is the absence of any belief on my part that the latter belief evidentially supports the former. But, intuitively, if I did hold such a belief, and it were operative in producing my

belief that my hands are shaking, the latter *would* be an inferred belief, despite the absence of a “recognizable inference pattern.”

Similarly, consider an individual who comes to believe that the moon is made of green cheese as a result of her belief that Socrates is mortal. *Could* the former belief be inferred from the latter? This depends on the reason why she engaged in that particular belief-state transition. If she is merely the victim of a strange disposition to believe that a celestial body is made of green cheese whenever she comes to believe that a major philosopher is mortal, then it is not a case of inference. If, however, she *believes*, for whatever reason, that the belief that Socrates is mortal evidentially supports the belief that the moon is made of green cheese, and, moreover, this belief (that the one supports the other) plays a causal role in her coming to believe that the moon is made of green cheese, then it is a case of inference.

By the same token, we can reconsider the case of the deranged individual who believes that the President is in the room whenever he believes that at least three people are in the room. If in fact, contrary to the manner in which I originally described the case, the individual takes the belief that the President is in the room to be evidentially supported by the belief that there are at least three people in the room, and this fact plays a causal role in his coming to believe that the President is in the room, then this *is* in fact a case of inferred belief. But if the crucial connecting belief is absent, namely, that the belief that there are at least three people in the room evidentially supports the belief that the President is in the room, then we do not have a case of inference.<sup>38</sup>

---

<sup>38</sup>Winters could try to argue that an inferential disposition which is triggered by a connecting belief *automatically fits* a recognizable pattern, for the connecting belief is then a premise of the inference, asserting that the other premise (or premises) of the inference evidentially support the conclusion inferred. “Surely,” she might say, “an inference to the conclusion that q, from beliefs that p, and that p evidentially supports q, fits a recognizable pattern.” I have two responses to this argument. First, as I discuss further below, I believe it a mistake to regard the connecting belief as a premise in the inference. Second, the suggested argument does not count *against* the view that connecting beliefs play a role in inference, but at most shows that the appeal to connecting beliefs does not

Robert Audi very nicely expresses the intuitive rationale for this connecting belief requirement. He says that when a person S infers a belief that p from another belief r,

he believes p in the light of r, not merely because of it... [H]e must in some way see r as supporting p. In part, the idea is that otherwise his belief that p is *cognitively unmotivated*: the reason S has for it, r, figures in the reason *why* he believes p, but his believing it is not *on account of* r, where this suggests, not his explicitly taking r into account, but his belief system's somehow reflecting r's subjectively registered support for p.<sup>39</sup>

So, inference is not merely a matter of being caused by one belief to hold another belief, or indeed of being disposed to form belief sequences of a certain general type. A reasoning subject is not a simple "input-output mechanism." A reasoning subject does not *simply* form a belief *in reaction to* other beliefs, but has a certain skill (be it innate or acquired) to form beliefs when she *detects*, or thinks she detects, that those beliefs are evidentially supported by certain other beliefs. In this sense, an inferred belief is inferred *in light of* the beliefs from which it was inferred. The deranged individual who falls prey to the belief-forming reflex of coming to believe that the President is in the room simply because of his belief that there are at least three people in the room does not believe the former in light of, or on account of the latter. He simply believes the former *because of* the latter, and *happens* to be disposed to repeat that process on other occasions. But the mere fact that he is disposed to repeat the process cannot make his belief an inferred belief.

Despite Audi's insight that connecting beliefs are required for inferences, his own account of inference is not quite correct. For he, like Armstrong in *A Materialist Theory*

---

preclude the possibility of explaining inference in terms of recognizable inference patterns. I have no quarrel with this last observation *per se*. I would only assert that an explanation of inference in terms of recognizable inference patterns is unhelpful in the absence of an appeal to connecting beliefs, and that given such an appeal, an explanation in terms of recognizable inference patterns is unnecessary.

<sup>39</sup>Audi (1993), pp. 240-241. Audi speaks, not of inference, but of "S believing p for a reason r." I assume that it is safe to reexpress this as S's inferring p from r, just as we did with respect to Armstrong's view above.

*of the Mind*, does not explain inferences in terms of belief-forming habits or dispositions. On his view, inference requires that the premise beliefs, together with the connecting belief, causally produce (and indeed sustain) the concluding belief, but he makes no specific mention of a requirement that the subject be disposed to make transitions of this sort on other occasions. As a result, his account is subject to certain difficulties, of which he himself is aware. He considers, for example, the following case:

Imagine that, in the presence of a machine which has the capacity to influence the brain so as to produce beliefs, S happens to form the beliefs that r, and that r warrants p.<sup>40</sup> It might also just happen that the machine is so adjusted that the formation of these beliefs causes it to produce in S a belief that p. The machine might at the same time suppress the normal causative processes by which the first two beliefs would produce the third, so that the third is not overdetermined, but generated by the first two only through the accidental activation of the machine.<sup>41</sup>

Audi (rightly) suggests that the belief that p is not *inferred* in this case despite the fact that the subject has a connecting belief. Audi's response to this and similar cases is to hold that inference requires a "non-wayward" causal connection between the premise beliefs (together with the connecting belief) and the conclusion.<sup>42</sup> He says also that the causal connection must be "non-accidental," and that it must not involve "alien intermediaries," that is, intermediaries, such as the belief-producing machine in the present example, which are "alien" to the subject drawing the conclusion.<sup>43</sup>

But I think that there is a much neater and simpler solution to the problem presented by this example. The case just discussed is not one of inference because the subject is not led to his belief that p, ( via his belief that r), because of *his disposition* to make transitions of that sort. Even if in fact the subject has such a disposition, it is not

---

<sup>40</sup>We may assume that the belief that r warrants p is an acceptable connecting belief, whatever 'warrants' might mean in this context.

<sup>41</sup>*Ibid.*, p. 254.

<sup>42</sup>*Ibid.*, p. 22.

<sup>43</sup>*Ibid.*, p. 262.

operative in the present case. Instead, the *machine* is ultimately responsible for his coming to believe that *p*.<sup>44</sup>

The problem with the disposition-based accounts discussed in the last section is that they did not appeal to connecting beliefs. At the same time, connecting beliefs by themselves are not sufficient to rule out problem cases. The solution is to combine the disposition-based account with the connecting belief requirement, yielding the following conception of inference:

- (I\*) S's belief that *p* is inferred from her beliefs  $B_1$  through  $B_n$  just in case:
- (i) S believes that her beliefs  $B_1$  through  $B_n$  evidentially support the belief that *p*.
  - (ii) The sequence of beliefs  $\langle B_1, \dots, B_n, p \rangle$  is of a kind *K* such that S has the following disposition *D*: If S has the first *n* beliefs in a sequence of kind *K*, and S holds an appropriate connecting belief that these first *n* beliefs evidentially support a further belief which would complete a sequence of kind *K*, and these beliefs are simultaneously activated, then S is disposed to form that further belief which would complete the sequence of kind *K*.
  - (iii) S's coming to believe that *p* is in fact a result of her disposition *D*, triggered by the simultaneous activation of her beliefs  $B_1$  through  $B_n$ , together with her belief that these *n* beliefs evidentially support the belief that *p*.

This account of inference raises a couple of issues which warrant comment.

First, the relationship, on this view, between the connecting belief in an inference and the premises of that inference is a potential source of confusion. If I believe that Jones would classify himself as a conservative because I believe that he is a member of the National Rifle Association and because I believe that ninety percent of the members of the National Rifle Association would classify themselves as conservative, the connecting belief required to make this transition an inference is *not* the latter belief. My belief that ninety percent of the members of the National Rifle Association would classify themselves as conservative is *not* the connecting belief which renders my belief-state

---

<sup>44</sup>Perhaps this is essentially what Audi means in saying that the causal connection must not involve alien intermediaries. If so, then I am merely offering an alternative way of expressing that requirement.

transition an inference. The connecting belief is a belief *over and above* the premises of the inference, and plays an entirely different role. The premises play an evidential role in supporting, or seeming to support, the conclusion. The connecting belief adds nothing to the support provided by the premises; its role is not evidential, but primarily explanatory. It reflects the fact that the subject *appreciates* the support given her conclusion by her premises, and that this appreciation is part of the reason why she reaches that conclusion.<sup>45</sup>

Second, there is a question about the required *content* of connecting beliefs. According to (I\*), an inference requires that the subject believe *that* her premises support her conclusion. But it seems reasonable that other sorts of connecting beliefs would serve just as well. For example, a subject might believe that her premises *adequately confirm* her conclusion, or that they make this conclusion *very probably true*. There are a multitude of propositional contents which connecting beliefs might have, but for simplicity I will subsume all of these under the rubric of believing that the premises of the inference evidentially support the conclusion.<sup>46</sup>

---

<sup>45</sup>There is a second reason why connecting beliefs should not be thought of as premises in inferences. If connecting beliefs were premises, then an infinite regress of connecting beliefs would be involved in any inference. For if the subject is required to believe that her beliefs B<sub>1</sub> through B<sub>n</sub> evidentially support her inferred belief that p, and this connecting belief is a premise in her inference, then if she is required to believe that her premises evidentially support her conclusion, she must believe that her beliefs B<sub>1</sub> through B<sub>n</sub> and that B<sub>1</sub> through B<sub>n</sub> evidentially support p in turn evidentially support p. But if this second connecting belief is likewise a premise in her inference, she will be required to have yet a third connecting belief, and so on. So, it is best not to think of connecting beliefs as premises.

<sup>46</sup>Hilary Kornblith suggested to me the possibility that inferences might be governed not by connecting beliefs but by "subdoxastic" awarenesses of evidential support relations--that is, awarenesses which either lack explicit propositional content or which have a more rudimentary content than is typically found in belief states. Whether such subdoxastic awarenesses present a viable theoretical alternative to connecting beliefs is a question which will have to await future research. It seems to me, however, that if such awarenesses are to play a role in guiding inference, they must be to some extent propositional, in which case they could still be considered connecting beliefs.

We have arrived at an account of the nature of inference, namely (I\*). But we have yet to discuss the conditions of *correct* inference. That problem is the focus of the next section.

### 5. *Blind and Mediated Tracking.*

When is an inferred belief *correctly* inferred? One obvious suggestion is that an inference is correct just in case it satisfies conditions (i) through (iii), set down in (I\*) above, *and* the premises of the inference do in fact evidentially support the conclusion. If an inference is such that the premises purport to provide evidence for the conclusion, then it is plausible to suppose that a correct inference is one in which the premises actually provide such evidence. There are counter-examples to this proposal, however.

Consider, for instance, this case due to Swain:

Suppose two subjects, Smith and Jones, who have the same evidence (grounds) for the belief that  $p$ , where the evidence consists of justified beliefs in the proposition  $p \vee (p \ \& \ q)$ . Both subjects come to believe that  $p$  on the basis of this evidence (and no other evidence). In the case of Smith, the mechanism for generating the belief is an inference which instantiates a tendency to invalidly infer  $p$  from any sentence of the form ' $p \vee q$ '. In the case of Jones, the mechanism is an inference which is based on an internalized valid inference schema (of which several are possible). It seems clear to me that only Jones has a justified belief that  $p$ , even though they have the same grounds.<sup>47</sup>

If we agree with Swain that Jones's belief is justified, while Smith's is not, this can only be because Jones's belief that  $p$  is correctly inferred, while Smith's belief is not. But, notice that, first, in both cases the premise evidentially supports the conclusion, and, moreover, conditions (i) through (iii) of (I\*) are plausibly satisfied. Conditions (ii) and (iii) are clearly satisfied insofar as both Smith and Jones arrive at a belief that  $p$  by way of general dispositions to make transitions of a particular sort. And it is not unreasonable to attribute appropriate connecting beliefs to Smith and Jones, albeit, perhaps, connecting beliefs which are operating "below the surface." that is, connecting beliefs which are not

---

<sup>47</sup>Swain (1988), pp. 468-469.

held consciously in mind (though they are, nevertheless, causally active). Therefore, a correct inference is not simply one which satisfies conditions (i) through (iii) of (I\*) and in which the premise (or premises) evidentially supports the conclusion, otherwise both Smith and Jones would have correctly inferred beliefs that p.

What then is wrong with Smith's inference process? According to Swain, the problem is that Smith is not a reliable information processing mechanism, at least in the situation in which he draws the specified inference. In determining whether a given inference is correct, one must, on Swain's view, take into account "relevant characteristics" of the subject, that is, characteristics which are probabilistically relevant to the reliability of the subject's information processing capacities. What is required for correct inference, on this conception, is that the subject, in forming a belief on the basis of certain premises, given his relevant characteristics as an information processor, be a reliable indicator of the truth of the inferred belief. Smith, in forming his belief that p, in the way that he does in the situation described, is not a reliable indicator of the truth of p, in that his belief that p is not very likely to be true, given the way in which he comes to hold it.<sup>48</sup>

I think that this assessment of Smith's inference process is essentially correct, but that it rests on a deeper explanation. The reason why Smith, in the situation in question, is a defective information processor is that while he draws a conclusion which is evidentially supported by his premise, he is not *sensitive* to the relevant relation of evidential support in drawing his conclusion; it plays no role in his cognition. Part of the problem can be put by saying that Smith's inference process does not "correspond" to the support relation in question in the sense that Smith could be led by his present inferential disposition, in a different situation, to form a belief which is not evidentially

---

<sup>48</sup>Ibid., p.469. See also Swain (1981), pp. 104-115.

supported by his premise.<sup>49</sup> In a situation in which he believes, say, 'p v r,' he will come to hold a belief which is not evidentially supported by his premise. Jones, on the other hand, will never be led by his inferential disposition to form a belief which is not evidentially supported by his premise. This is one of the crucial reasons why Jones's belief is correctly inferred, while Smith's is not.

Let us explore this further. In the preceding two sections, we spoke of inferred beliefs as beliefs which are generated by a disposition to form sequences of beliefs of a certain kind K. Now, it is plausible to assume that when a set of n beliefs evidentially support some other belief, then these beliefs taken together form a sequence of length n+1 which is of a general kind such that the first n beliefs of a sequence of that kind always evidentially support the n+1<sup>th</sup> belief. In that case, we can imagine a correspondence between certain belief-forming dispositions, and certain relations of evidential support. Such a correspondence exists when one is disposed to form belief sequences of a general kind K, and every sequence of kind K is such that its first members evidentially support its final member.

So, correct reasoning involves, in part, a correspondence between a subject's inferential disposition and a relation of evidential support. But, as I argued in the preceding sections, inference also requires a connecting belief which is operative in the transition from premises to conclusion. In a correct inference, a connecting belief guides the inferential transition, so that the inferential transition follows an evidential support relation in virtue of the connecting belief. In this sense, then, a correct inference does not simply correspond to a relation of evidential support, but, if you will, *tracks* that support relation, in that the correspondence between the inferential disposition and the support

---

<sup>49</sup>Firth (1978) suggests that a belief is correctly inferred from its premises (or based rationally on its premises, as he says) if there is a "correspondence" between the psychological processes by which that belief is generated, and the evidential relations by which it receives support from its premises. I follow his suggestion here.

relation does not occur by accident, but because of an appropriate connecting belief held by the subject. Thus, I suggest that a correctly inferred belief is a belief produced by an inferential disposition which tracks a relation of evidential support (under the guidance of a connecting belief), a proposal which we may call, for brevity, the tracking account of correct inference.<sup>50</sup> The following is a more precise statement of this view:

- (CI) S's belief that p is correctly inferred from her beliefs B<sub>1</sub> through B<sub>n</sub> just in case:
- (i) The sequence of beliefs <B<sub>1</sub>, ..., B<sub>n</sub>, p> is of a kind K such that the first n members of the sequence evidentially support the final member,
  - (ii) S believes that her beliefs B<sub>1</sub> through B<sub>n</sub> evidentially support her belief that p,
  - (iii) S has the following disposition D: If S has the first n beliefs of a sequence of kind K, and S believes that these n beliefs evidentially support some further belief which would complete a sequence of kind K, and these beliefs are simultaneously activated, then S is disposed to form the belief which complete the sequence of kind K, and,
  - (iv) S's coming to believe that p is in fact a manifestation of her disposition D, triggered by the simultaneous activation of her beliefs B<sub>1</sub> through B<sub>n</sub>, together with her belief that these n beliefs evidentially support the belief that p.

Relative to this account, we can explain why Jones's belief is correctly inferred while Smith's belief is not. Jones's belief satisfies each of the four conditions in (CI). The sequence of beliefs < p v ( p & q ), p > is of a kind K of belief sequences such that the first member of a sequence of that kind evidentially supports its final member. (Kind K, in this case, presumably consists of 2-membered sequences, the first member of

---

<sup>50</sup>More specifically, as I shall explain below, this is the *mediated* tracking account of correct inference, as opposed to what I call the *blind* tracking account. I shall explain the distinction between these two views shortly. The idea of a tracking account of correct inference was inspired by Nozick's (1981) account of knowledge as belief which tracks the truth, in the sense that if the belief were not true, the subject would not hold it, and if it were true, the subject would hold it. It occurred to me that a correctly inferred belief might be one which tracks, not truth, but evidential support, in the sense that if the belief were not evidentially supported, the subject would not hold it, and if it were evidentially supported, the subject would hold it. But, as a result of the issues discussed in this chapter, the account of correct inference turns out to be quite a bit more complicated than this.

which is a disjunction, and the second member of which is evidentially supported by both disjuncts of the first member.) The second condition, as I indicated above, is satisfied, insofar as Jones' belief that  $p$  is inferred, and inference requires the presence of a connecting belief. Finally, and crucially, Jones is disposed to form belief sequences of the very kind in question, that is, when he holds a disjunctive belief, both disjuncts of which evidentially support some further belief, he is disposed to form that further belief, provided he holds an appropriate connecting belief, and it is this inferential disposition which is operative on the occasion in question. Thus, he satisfies each of the conditions (i) through (iv) in (CI), and his inference is thereby correct.

Smith, on the other hand, fails to satisfy condition (iii). That is, Smith is *not* such that when he holds a disjunctive belief, both disjuncts of which evidentially support some further belief, he is disposed to form that further belief. His disposition leads him to believe the first disjunct of disjunctive beliefs that he holds, regardless of whether the latter evidentially supports the former. The problem, then, is that his belief sequence  $\langle p \vee (p \& q), p \rangle$  on the occasion in question belongs to two kinds. On the one hand, it is of a kind  $K$  such that its first member evidentially supports its final member, thus satisfying condition (i). On the other hand, it is of a kind  $K'$  such that the second member is one of the disjuncts of the first member. It is belief sequences of this second kind  $K'$  which Smith's inferential disposition leads him to form, not belief sequences of kind  $K$ , thus violating condition (iii).<sup>51</sup>

Despite its success in explaining the incorrectness of Smith's inference, some theorists will worry that the tracking account of correct inference contained in (CI) imposes too stringent a requirement on cognitive subjects. If a subject is to satisfy

---

<sup>51</sup>We could describe the case so that Smith *is* disposed on certain occasions to form a belief which is evidentially supported by both disjuncts of a disjunctive belief that he holds. It would then be condition (iv) which is violated, since it is not *this* disposition which is operative on the occasion we are considering.

condition (ii), which we may call the connecting belief requirement, then she must possess certain evidential concepts, in order that she be able to form beliefs about the evidential support which her premises provide for her conclusion. Although, as I discussed in the last section, the subject need not believe *that* her premise beliefs evidentially support her conclusion, she must hold an at least analogous belief, asserting some sort of evidential relationship between her premises and her conclusion (e.g., that the premises adequately confirm the conclusion, or make it very probably true). While it is clear that most ordinary adult human beings operate with at least some evidential concepts, and thus have the capacity to form connecting beliefs, it is equally clear that many beings (such as small children and animals) who appear to engage in processes which are at least similar to adult human reasoning do not operate with such concepts. Thus, to take a natural example, when a dog becomes excited upon seeing his leash, we could say that he undergoes a process which is at least *like* reasoning in that he forms an expectation or "belief" that he is about to go for a walk as a result of his "belief" that his master has his leash. But the dog seems incapable of forming connecting beliefs, at least if having such beliefs would require the dog to possess concepts of evidential support. According to (CI), then, the dog is not, strictly speaking, reasoning in coming to believe that he is about to go for a walk.

Some theorists will find this fact about (CI) objectionable. For example, Barbara Winters, who rejects the connecting belief requirement, claims that one of the virtues of her account is that it "enables us to see human pathways to belief as not different in kind from the rest of the animal kingdom."<sup>52</sup> If the intuition is quite strong (and widely shared) that children and relatively intelligent animals reason *in the same sense* that adult human beings do, then condition (ii) in (CI) must either be rejected, or altered in some way so that animals and small children more easily satisfy it. One possibility would be

---

<sup>52</sup>Winters (1983), pp. 217-218.

to follow Robert Audi in allowing *de re* beliefs to fill the role of connecting beliefs.<sup>53</sup> On Audi's view, a subject need not believe *that* her premises evidentially support her conclusion: she is merely required to believe *of* a support relation that it holds between her premises and her conclusion, and thus need not subsume the support relation under any particular concept or description. To illustrate, the support relation in a specific case might be one of entailment, but the subject is not in Audi's opinion required to believe *that* her premises entail her conclusion. She must only believe *of* entailment that it holds between her premises and her conclusion. On this *de re* conception of connecting beliefs, subjects who lack concepts of evidential support may still be capable of forming connecting beliefs, and may well satisfy appropriately modified versions of the conditions in (CI). As Tomis Kapitan, who accepts Audi's suggestion that *de re* beliefs can play the role of connecting beliefs, explains:

The one-year-old who turns her head towards the door upon hearing her father's voice in the next room exhibits behavior suggesting awareness of a connection between the occurrence of a certain type of sound and the presence of a particular person. The same can be said of the dog who runs to the door upon seeing his master with the chain. Of course, *we* might report the contents of their awareness with terms like 'implies' or 'makes likely', but these can be construed as having only an *external* <sup>54</sup>occurrence in such contexts.... What is crucial is that *some* sort of connection is grasped, and this need not entail that the agent has the concepts of deductive or inductive *relations* among propositions....<sup>55</sup>

The problem with this suggestion, however, is that it is doubtful that animals and small children have the conceptual resources to form even *de re* connecting beliefs, for such beliefs still involve propositional content which can be comprehended only by

---

<sup>53</sup>Audi (1993), p.241. The following is an example of a *de re* belief. Imagine that the President is dressed as a clown as part of a clever plan to disguise and protect him. I might see him dressed in this outfit and, not recognizing him, form the belief that the man before me is a buffoon. Then I would believe *of* the President that he is a buffoon (a belief *de re*), though I would not believe *that* the President is a buffoon.

<sup>54</sup>Kapitan means, in other words, that we can attribute to small children and animals *de re* beliefs about evidential support relations, without implying that these agents have concepts of evidential support, or at any rate, *our* concepts of evidential support.

<sup>55</sup>Kapitan (1985), p. 282.

relatively sophisticated beings (the subject believes of a support relation *that* it holds between certain beliefs). Thus, the claim that animals and small children form beliefs under the guidance of *de re* connecting beliefs would seem to be motivated only by the intuition that such creatures reason in the same sense in which adult human beings do. How strong is this intuition? Could not we just as well say that what animals and children do is an activity like adult reasoning in some respects, but unlike it in others. We might say that animals and small children engage in *proto*-reasoning, which, in the case of children, gives way in the fullness of time to genuine reasoning. This *proto*-reasoning is like fully developed adult reasoning in that it involves a disposition to form beliefs when prompted by the appropriate prior beliefs.<sup>56</sup> We could even define a sense in which such *proto*-reasoning tracks relations of evidential support. For it could be that dogs and other animals have over time evolved certain *proto*-inferential dispositions precisely because those dispositions correspond to relations of evidential support.<sup>57</sup> We might refer to this as *blind tracking*, because it occurs without the guidance of connecting beliefs. Accordingly, we could refer to the account contained in (CI) as the *mediated tracking* account, since it requires the tracking relationship between disposition and evidential support to be mediated by an appropriate connecting belief. The *blind tracking* account, in contrast to (CI), would look like this:

- (BT) S's belief that p is produced by correct *proto*-reasoning from her beliefs B<sub>1</sub> through B<sub>n</sub> just in case:
- (i) The sequence of beliefs <B<sub>1</sub>, ..., B<sub>n</sub>, p> is of a kind K such that the first n beliefs evidentially support the final belief,
  - (ii) S has the following disposition D: If S has the first n beliefs in a sequence of kind K, then S is disposed to form a further belief which would complete a sequence of kind K,

---

<sup>56</sup>One question, of course, is whether animals and small children have beliefs in the same sense as fully developed adult human beings. One possibility is that such beings have *proto*-beliefs, i.e., belief-like states, but with a cruder or more rudimentary propositional/conceptual content than the belief states typically found in adult human beings.

<sup>57</sup>Nozick (1993) has made this suggestion, as I shall discuss in Chapter 5.

- (iii) S's coming to believe that p is in fact a manifestation of her disposition D, triggered by her beliefs B<sub>1</sub> through B<sub>n</sub>.

Of course, the introduction of this second tracking notion raises the following obvious question: If it is possible for (proto-)reasoning processes to track relations of evidential support in the absence of connecting beliefs, why postulate connecting beliefs, even in the case of ordinary adult human beings? After all, one could argue that even if most ordinary adults have the conceptual capacity to form connecting beliefs, such beliefs are comparatively rare. Is it really plausible to think, one might ask, that Jones, in the case to which we have frequently returned, is guided by a connecting belief that his belief that a can of lighter fluid is very near an open flame evidentially supports his belief that he is in a dangerous situation? Is the connecting belief in fact required to explain the transition? Perhaps Jones' belief that he is in a dangerous situation is produced by proto-reasoning in the sense of (BT). But in that case, one could object, proto-reasoning just *is* reasoning, and all we need for an adequate theory of correct human inference is (BT), not (CI) with its superfluous connecting belief requirement.

Let me make four points in defense of (CI). First, one explanation for the sense that Jones is not guided by a connecting belief in coming to believe that he is in a dangerous situation is that the connecting belief requirement is, perhaps, easily confused with the intellectualist model of reasoning. But insofar as connecting beliefs can be causally operative in guiding reasoning without being consciously before the mind of the reasoning subject, the connecting belief requirement does not render reasoning overly deliberate or rehearsed. From this point of view, then, there is no reason not to think that connecting beliefs are involved in cases in which reasoning occurs in an automatic or habituated way. Of course, it may be that if one explicitly asked a subject in such a case whether she were relying on a connecting belief in drawing some inference, she might answer that she was not aware of any such belief. But then again, the subject might also

claim not to be aware of certain *premises* which were clearly causally active in her inference, but which were not consciously before her mind.

Second, as I have tried to argue above, connecting beliefs play a role in helping us to make sense out of reasoning on a philosophical level, in determining the features which separate reasoning from other sorts of belief generating processes. The central idea is that an inferential transition is rationally intelligible (it *makes sense*) because the subject in reaching her conclusion purports to arrive at a belief which is evidentially supported by her premises. If the connecting belief explanation of the rationally intelligible character of inference is rejected, then some other explanation must be given. The only other possibility on offer is Winters's "recognizable inference pattern" requirement, which we have seen is unsatisfactory. Until a better explanation becomes available, then, the connecting belief requirement offers the best means of distinguishing between inferences and noninferential habituated belief-state transitions.

Third, we can easily imagine cases in which connecting beliefs would help to explain why a subject comes to hold a certain belief, as opposed to some other belief. If Jones infers a belief that  $p$  from a belief that  $(p \ \& \ r) \vee (p \ \& \ q \ \& \ r)$ , but does not infer the belief that  $r$ , one might best explain this by saying that Jones noticed (i. e., believed) that  $p$  was supported by his premise, but did not notice that  $r$  was supported by his premise. If one attempts to explain the transition simply by appeal to a blind disposition, say, that Jones tends to infer a belief when it is evidentially supported by both disjuncts of a disjunctive belief that he holds, then one may not be able to explain why Jones came to believe that  $p$  and not that  $r$ .

Finally, fourth, we must keep in mind that a correctly inferred belief should contribute to a subject's being in a secure position with respect to the truth. From this perspective, (CI) is better than (BT) as an account of correct inference, since the secureness of a subject's position in reaching some conclusion would seem to be

undermined if she were not guided by an appropriate connecting belief. For example, the secureness of Jones's position in reaching his conclusion that he is in a dangerous situation seems to depend on his having some awareness of the fact *that* the can of lighter fluid's being very near an open flame is evidence for his conclusion. Furthermore, we can explain why certain cognitive subjects (specifically adults who operate with evidential concepts) should want to satisfy the more stringent standards of (CI), since it is these subjects who are most appropriately evaluated with respect to the secureness of their positions in holding the beliefs they do. It is less appropriate to subject animals and small children to this sort of evaluation, which explains why (BT) might be thought a more useful tool for assessing their cognitions.

But, for all of these points in favor of (CI), it may still be that much of human reasoning is "blind" (or, to put this another way, that much of what we take to be reasoning in humans is in fact proto-reasoning). After all, the question whether a given belief-state transition occurs under the guidance of a connecting belief is, in the final analysis, an empirical question, to be answered by appropriate psychological investigation of the reasoning subject. It may be that such investigation would reveal in many instances the absence of a connecting belief. In that case, it might be more plausible to regard (CI) not as an account of correct inference as it usually occurs, but rather as the expression of an ideal to which we would *like* our inferences conform, and to which they would conform were it not for our cognitive limitations.

Before concluding this section, I want to return to the rule based account of reasoning discussed earlier in the chapter. The advocate of the rule-based account might make the charge that the account I have defended here can be translated into the terms of the rule-based account. Thus, whereas I speak of a subject coming to hold a belief that *p* as the result of a disposition which tracks a relation of evidential support under the guidance of a connecting belief, the rule-based advocate could claim that the subject's

reasoning is governed by the following rule: from beliefs  $B_1$  through  $B_n$  and the belief that  $B_1$  through  $B_n$  evidentially support  $p$ , one is permitted to infer a belief that  $p$ . The rule-based advocate could claim that the subject's belief state-transition is explained by the fact that it conforms to this rule, that the subject undergoes the transition because it satisfies the rule.<sup>58</sup>

One problem I have with this rule-based way of rephrasing my theory is that it gives the impression that the connecting belief is a *premise* in the inferential transition, insofar as the connecting belief is specified in the antecedent of the rule. Aside from this point, however, I have no difficulty *per se* with this rule-based attempt to reformulate my account. I would only insist that the disposition-based approach gets at a deeper level of explanation than the rule-based account, that what ultimately explains both the nature of reasoning, and the nature of correct reasoning, are the inferential dispositions of cognitive subjects, and features thereof.

One task remains to be completed, namely, converting the account of correct inference offered here into an account of correct inferential basing.

#### 6. *The Basis Condition Restated.*

I indicated in the first section of this chapter that I would make the idealizing assumption that an individual holds a belief on the inferential basis of certain evidentiary beliefs only if she inferred the former from the latter at some previous time. Thus, at a particular time  $t$ , a subject  $S$ 's belief that  $p$  is correctly inferentially based on her beliefs  $B_1$  through  $B_n$  if there was a previous time  $t'$  at which she correctly inferred her belief that  $p$  from her beliefs  $B_1$  through  $B_n$  and in the intervening time, her belief that  $p$  has been causally sustained by the latter beliefs. This leads to the following reformulation of the Basis Condition:

---

<sup>58</sup>such a suggestion is made in Talbott (1990b), p. 63.

- (BC\*) S's belief that p is inferentially justified, at a time t, by her beliefs B<sub>1</sub> through B<sub>n</sub> only if:
- (i) The sequence of beliefs <B<sub>1</sub>, ..., B<sub>n</sub>, p> is of a kind K such that the first n members evidentially support the final member.
  - (ii) S believes that her beliefs B<sub>1</sub> through B<sub>n</sub> evidentially support her belief that p.
  - (iii) S has the following disposition D: If S has the first n beliefs in a sequence of kind K, and S believes that these n beliefs evidentially support some further belief which would complete a sequence of kind K, and these beliefs are simultaneously activated, then S is disposed to form the further belief which would complete the sequence of kind K.
  - (iv) There was a time t', not later than t, such that S held at t' the beliefs B<sub>1</sub> through B<sub>n</sub>, and believed that these n beliefs evidentially support the belief that p, and S came to believe that p at t' as a result of her disposition D, triggered by simultaneous activation of her beliefs B<sub>1</sub> through B<sub>n</sub> together with her belief that these n beliefs evidentially support the belief that p, and.
  - (v) For any time t'', no earlier than t' and no later than t, S's belief that p is sustained at t'' by her beliefs B<sub>1</sub> through B<sub>n</sub>, together with her belief that these n beliefs evidentially support her belief that p.

In view of my idealizing assumption, this new version of the Basis Condition ignores the psychological reality that individuals often come to hold beliefs on the inferential basis of evidentiary beliefs from which they were not originally inferred. For example, an employer might initially believe that a certain applicant is the best candidate for a job for entirely superficial reasons, such as the applicant's taste in literature or the fact that she was born on the East Coast. But over time, these superficial reasons might come to be wholly replaced by legitimate reasons concerning the applicants work history or education. The crucial feature of such cases seems to be that while the subject did not infer the belief in question from the evidentiary beliefs on which it comes to be based, she could (or would) have if conditions had been different. But the issue is complicated, and I can offer no detailed proposal to deal with these sorts of cases at this time.

## 7. Summary.

In this chapter, I discussed the nature of inference, and advocated a disposition-based account of that notion, as against a rule-based account. That is to say, I took the view that a given belief-state transition qualifies as an inference only if that transition is a manifestation of a disposition on the part of the subject to make transitions of a particular sort; the question whether the belief-state transition is an inference is not directly concerned with whether it is “governed by rules.” I also advocated, more specifically, that inferences must be guided by *connecting beliefs*, that in drawing an inference a subject is prompted by a belief which she possesses concerning a relation of evidential support between her premises and her conclusion.

The main purpose of this chapter has been to work toward a more complete and detailed statement of the Basis Condition which was introduced in Chapter 2. This reformulation of the Basis Condition was derived from a conception of correct inference as a process which *tracks* evidential support relations via the connecting beliefs which I have argued are involved in our inferences. But, recognizing the controversial character of this mediated tracking account, as I called it, I sketched alongside it a *blind tracking* account which should appeal to those advocates of the disposition-based approach who reject the connecting belief requirement on inference.

## Four

### The Access Condition

#### 1. Overview.

In the first chapter, I argued that only internalist theories of inferential justification can adequately capture the cognitive desideratum of being in a secure position with respect to the truth. The distinctive mark of such theories is the acceptance of what may be called an "access condition" specifying certain justificatory factors which must be reflectively accessible to a subject if she is to have an inferentially justified belief. In Chapter 2, I gave a rough formulation of such an access condition, which I shall try to improve upon in the present chapter. My concern is thus to answer the following two questions. First, what justificatory factors must be accessible on reflection if a person is to have an inferentially justified belief? Second, what does it mean to say that these factors are reflectively *accessible*--that is, that the subject *can* access them? These questions will be examined in sections three and four respectively. First, however, I must explain why the preliminary formulation of the Access Condition arrived at in Chapter 2 is unacceptable as it stands.

The Access Condition was initially stated in the following way:

- (AC) S's belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if she can come to know on reflection that her belief that p is correctly inferentially based on her beliefs B<sub>1</sub> through B<sub>n</sub>.

As I will discuss in the next section, this formulation is unacceptable because a belief's being correctly inferentially based on its evidence is not something which a subject can *noninferentially* know. On pain of a vicious regress, only justificatory factors which can be noninferentially known to obtain may be required to be reflectively accessible. The object, then, will be to replace the present formulation of (AC) with one which captures

the spirit of that formulation while requiring only noninferentially knowable justificatory factors to be reflectively accessible.

In the search for an acceptable formulation of (AC), it is helpful to return to a distinction among internalist theories, due to Robert Audi, which we encountered briefly in the second chapter. According to this distinction, there are on the one hand theories which require only that the subject have access to *what justifies* a given belief, which Audi calls *first-order* internalist theories. On the other hand, there are theories which require access not only to what justifies the belief, but also at least some access to *how* what justifies does so. These he calls *second-order* theories.<sup>1</sup> First order theories have been advocated by prominent internalists, including Audi himself. The basic problem with such theories, as I shall argue, is that they fail to make sense as accounts of what it is to be in a secure position with respect to the truth. To make sense as an account of what it is to be in a secure position with respect to the truth, a theory must require access not only to what justifies a given belief (i.e., the particular evidentiary beliefs which confer justification on it), but also to at least some aspects of how what justifies does so. As it presently stands, (AC) represents a form of second-order internalism, since it requires a subject to have access not only to the fact that she holds certain evidentiary beliefs, but also to the fact that her belief is correctly inferentially based on them. If the current incarnation of this condition is unacceptable, then the natural course is to isolate certain noninferentially knowable factors which go into a beliefs being correctly inferentially based on its evidence, and require *these* to be reflectively accessible. In this way, hopefully, an acceptable second-order internalism which approximates the original (AC) without inheriting its defects can be achieved.

---

<sup>1</sup>Larry BonJour has pointed out to me that this terminology is not quite apt since the relations whereby a set of evidentiary beliefs confer justification on some further belief are part of *what justifies* that further belief. The distinction is nevertheless convenient, and I will continue to use it.

The version of second-order internalism for which I shall argue requires that a subject have access on reflection to relations of evidential support and causal sustenance which hold between her belief and its evidence. I shall try to show that this preserves some measure of the requirement that the subject have access to *how* her evidence justifies her belief, while requiring only noninferentially knowable justificatory factors to be reflectively accessible. Of course, two objections will have to be faced: first, that these relations are *not* indeed reflectively accessible; and second, that a version of internalism which requires only these factors to be reflectively accessible is not plausible as an account of what it is to be in a secure position with respect to the truth. I will discuss the latter objection in section three below, but much of the discussion of the former will have to be postponed until Chapter 5.

## 2. *Reflective Access and Noninferential Knowledge.*

So far, we have kept all talk of “reflective access” on a fairly intuitive level. We have said that, roughly speaking, a person has reflective access to the truth of a proposition *p* when she can come to know that *p* is true simply by carefully focusing her attention on the question of whether *p* is true. Standard examples of truths which can be known in this way include simple truths about the person’s present state of mind, as well as certain truths of logic and mathematics, and others traditionally recognized as knowable *a priori*. This is not meant to imply that there is no important distinction between introspective knowledge and *a priori* knowledge. The point is merely that for purposes of our discussion, we may regard these two sorts of knowledge as species of a general variety of knowledge we may call *reflective* knowledge.

The sorts of paradigm cases of reflective knowledge which we have discussed (e.g., the knowledge that I am currently imagining a purple beach ball, or that nothing can be red and green all over at the same time) are plausible examples of *direct* or

*noninferential* knowledge. When I reflect carefully on whether something could be red and green all over at the same time, and see thereby that no such thing could possibly exist, I do not infer this knowledge from any further premises; my knowledge appears to be quite direct and immediate. Similarly, when I come to know by focusing my attention that I am imagining a purple beach ball, I seem to know this without benefit of supporting premises, by directly apprehending the contents of my imagination. It is clear, however, that it is also possible to have *inferential* reflective knowledge. For example, if I conclude that I do not like detective novels after mulling over in my mind the fact that I have never been able to finish reading one, this belief seems to result through a kind of inferential reflection, in that I first bring to mind by focusing my attention the fact that I have never been able to finish reading a detective novel, and I then infer from this that I must not like them very much. Similarly, if I construct a proof in number theory that the square of an odd number is always odd, then it seems reasonable to say that I have *inferential* reflective knowledge that the square of an odd number is always odd. We might say, roughly, that I know something via inferential reflection if I have correctly inferred it from things that I reflectively know.

Now, it appears that if one is to know on reflection that a given belief is correctly inferentially based on its evidence, then this must be reflective knowledge of the inferential sort. For in the typical case at least, a belief is inferentially based on certain other beliefs at a time  $t$  only if it was at some earlier time  $t'$  inferred from these beliefs, and has been continuously causally sustained by them through the intervening time. But while one may (as I shall argue) be capable of knowing directly that a given belief is *now* causally sustained by certain other beliefs, and one may remember (and thus directly know<sup>2</sup>) that one inferred this belief from its evidence at some earlier time, one still seems

---

<sup>2</sup>Some epistemologists have argued that a belief about a past event may be directly justified by being appropriately based on a "memory impression" of that past event having occurred. See, for example, Audi (1993), p. 309, and (1988), ch. 2.

to be in the position of having to *infer* that one's belief has been continuously sustained by its evidence throughout the intervening time. Even if for any of the intervening moments one chooses to consider, one remembers that one's belief was sustained by its evidence at that time, one still must infer from this that the belief has been continuously sustained throughout the time in question.

This fact, that a belief's being correctly inferentially based on its evidence can only be known inferentially, raises a serious difficulty for the Access Condition as presently formulated. For suppose that some subject S has an inferentially justified belief that p. Then, by the Access Condition:

- (1) S can know on reflection that her belief that p is correctly inferentially based on its evidence.

But since S could only know this inferentially, her being able to know that her belief that p is correctly inferentially based on its evidence requires that she be able to satisfy the Access Condition with respect to this latter knowledge as well (for, of course, her inferentially knowing that p is correctly inferentially based on its evidence requires that she be inferentially justified in believing this). Thus, it must be the case that were S to form on reflection the belief, call it p', that her belief that p is correctly inferentially based on its evidence, she would be able to know on reflection that this second belief is based on *its* evidence. Thus, (1) is true only if:

- (2) S can know on reflection that her belief that p' (i.e., that her belief that p is correctly inferentially based on its evidence) is correctly inferentially based on its evidence.

Once again, the potential knowledge expressed in (2) would have to be inferential, requiring once more that S be able to satisfy the Access Condition. Thus, if S were to form on reflection the belief, call it p'', that her belief that p' is correctly inferentially based on its evidence, she would have to be able to tell on reflection that this third belief is

correctly inferentially based on its evidence. But since  $p'$  is the proposition that her belief that  $p$  is correctly inferentially based on her evidence, this requires that:

- (3) S can know on reflection that her belief that  $p''$  (i.e., that her belief that her belief that  $p$  is correctly inferentially based on its evidence is correctly inferentially based on its evidence) is correctly inferentially based on its evidence.

The knowledge expressed in (3) would *also* have to be inferential, and the cycle continues. The viciousness of this regress lies not in requiring a finite knower to have an infinite number of explicit items of knowledge (since at each stage only *potential* knowledge is required). Rather, the problem is that at each stage of the regress, the subject is required to be *able* to form a belief of greater propositional complexity than beliefs at earlier stages of the regress. At some point, a threshold will be passed in what the subject can plausibly be thought to comprehend, let alone determine on reflection.<sup>3</sup>

This regress shows that while reflective knowledge need not always be direct or noninferential, the justificatory factors required to be reflectively accessible by the Access Condition must be *capable* of being known to obtain noninferentially, otherwise, a subject would always be required to be able to satisfy the Access Condition with respect to the reflective knowledge the Access Condition itself requires her to have.<sup>4</sup> Thus, the present formulation of (AC) is unacceptable, and we must face the problem of replacing what it *does* require subjects to reflectively determine with justificatory factors that are

---

<sup>3</sup>This regress is similar to the one which besets the pure internalist position discussed in Chapter 2. Both regresses are generated by the fact that a subject in reflectively accessing some justificatory factor is required to have reflective access to some further more complicated justificatory factor, resulting in a hierarchy of ever increasing iterated that-clauses. There is an important difference between the two regresses, however. Pure internalism generates a regress by requiring a subject to tell on reflection *that* she can tell on reflection that some justificatory factor is present. The Access Condition, however, does not require a subject to tell on reflection that she can tell on reflection that her belief is correctly inferentially based on her evidence. Rather, it leads to a regress because a subject could only know inferentially that her belief is correctly inferentially based on its evidence, requiring her to satisfy the Access Condition with respect to this knowledge as well.

<sup>4</sup>It shows, moreover, that noninferential justification must not require satisfaction of an access condition, since otherwise an analogous regress would ensue.

noninferentially accessible. This reformulation must, however, be accomplished without losing the spirit of the condition's initial formulation, since that grew out of an effort to capture in a precise way the cognitive desideratum of being in a secure position with respect to the truth. The remainder of this chapter will focus on this task.

### 3. *First and Second Order Internalism.*

The trouble with (AC), as we have seen, is that it requires a subject to come to know on reflection that her belief is correctly inferentially based on others, something which could never be noninferentially known. Some theorists might diagnose the problem by saying that (AC) makes the mistake of requiring a subject to have access not just to *what* justifies her belief (i.e., the particular evidentiary beliefs on which it is based) but also to *how* what justifies does so. They might suggest that if the Access Condition is to be restricted so as to require access only to noninferentially knowable justificatory factors, then this condition can at most require a subject to have access to her evidentiary beliefs, not to any of the relations whereby these beliefs confer justification.

Several prominent internalists have adopted positions of this sort. For example, Robert Audi insists that "internalism requires that one have internal access (roughly, access by reflection) to *what* justifies, to justifying grounds; it need not require such access to *how* what justifies does so, e. g. (in part) by sustaining the justified belief (or even to the fact that it justifies)."<sup>5</sup> And he says elsewhere that "internalism does not imply that one can become aware of [one's justifying grounds] under any particular description, such as 'my evidence,'" further reinforcing the idea that a subject need only have access to her grounds, and not to the role that those grounds play in justifying her belief.<sup>6</sup> Likewise, Earl Conee, suggests that what can be "external to the mind of a

---

<sup>5</sup>Audi (1993), p. 30.

<sup>6</sup>Ibid., p. 332.

person whose belief is justified, i. e., inaccessible to the person by reflection, are epistemological facts about what evidence provides the person's justification and about the nature of the epistemic link of the belief to its justifying evidence. What must be internal, i.e., accessible to the person by reflection, is evidence that does in fact suffice to justify the belief."<sup>7</sup> Alston holds a third view along these same lines. He points out that, with regard to requirements of reflective access, one must distinguish between "(1) the 'justifier,' that is, the ground of the belief, and (2) its adequacy or justificatory efficacy: its truth conduciveness."<sup>8</sup> He accepts that the justifier or ground must be "the sort of thing whose instances are fairly directly accessible to their subject on reflection."<sup>9</sup> but he goes on to argue that while

it may well be within the limits of human capacity, it is by no means always the case that the subject of a justified belief is capable of determining the adequacy of his ground, just by careful reflection on the matter, or, indeed, in any other way. For one thing, many subjects are not at the level of conceptual sophistication to even raise the question of adequacy of ground, much less determine an answer by reflection. One thinks here of small children and, I fear, many adults as well.<sup>10</sup>

So, one reaction to the result arrived at in the last section, that only noninferentially knowable justificatory factors may be required to be reflectively accessible, is to require only that a subject have access to what justifies her belief, i.e., the particular evidentiary beliefs which confer justification on it, as opposed to the relations between her belief and its evidence on which this conferring of justification depends. As we have seen, Audi calls this position "first-order internalism," and contrasts it with "second-order internalism" which does require at least some access to the means by which a subject's evidence confers justification on her belief. Since first-order internalism does not require any access to the role which a subject's evidentiary beliefs

---

<sup>7</sup>Conee (1988), p. 400.

<sup>8</sup>Alston (1989), p. 237.

<sup>9</sup>Ibid., p. 238.

<sup>10</sup>Ibid., p, 240.

play in rendering her belief justified, the most that this view can require is that the subject have access to the *existence* of the evidentiary beliefs in question. The suggestion, then, is that we replace (AC) with:

(AC1) S's belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if S can know on reflection that she holds the beliefs B<sub>1</sub> through B<sub>n</sub>.

It seems to me, however, that this very minimal access condition on inferential justification, in the absence of any stronger access requirements, fails to make sense as part of an account of what it means to be in a secure position with respect to the truth. In order to plausibly explicate this notion, *some* access to how the evidentiary beliefs confer justification is necessary. Let me explain.

First, it makes little sense (in the context of an attempt to elucidate the notion of the secure position) to hold, as Audi does, that, on the one hand, justification requires that the grounds or evidence for a belief be accessible to the subject, but that, on the other hand, the subject need not access the grounds under any particular description, such as "my evidence." It is difficult to see how a subject's reflective access to the grounds for a certain belief could play any role in that belief's being justified for her, if she did not, in reflectively accessing those grounds, *appreciate* that they serve as her evidence. For, suppose that some subject's belief that, say, Bill Clinton will be in Seattle tomorrow is correctly inferentially based on her belief that his planned arrival was announced this evening on television, and that while she has reflective access to the fact that she believes the latter, she has no access to the fact that this belief serves as her evidence for believing that Bill Clinton will be in Seattle tomorrow. If someone were to ask her what reason she has for believing that Bill Clinton will be in Seattle tomorrow, what would she say? It is true that she has reflective access to the belief which could provide a proper answer to this question, but then again, she has reflective access to a multitude of other beliefs. What is to lead her to cite in response her belief that Bill Clinton's planned arrival was announced on television tonight, rather than, say, her belief that people once walked on the Moon, if

she has no appreciation of the fact that former belief (not the latter) is her evidence for believing that Bill Clinton will be in Seattle tomorrow? It seems, therefore, that a first-order internalist must at least concede that having an inferentially justified belief requires the subject to have access on reflection not simply to the grounds or evidence for that belief, but to the fact that those grounds *are* her evidence.

One way to spell this requirement out (that the subject have access to the fact that the grounds of her belief *are* her evidence) would be to require that the subject be able to tell on reflection that her belief is inferentially based on its grounds. We have seen, however, that this is not something which can be required to be reflectively accessible. A plausible alternative, on my view, is to require that the subject be able to tell on reflection that she holds her belief *because* she takes her grounds to provide good evidence for that belief, i.e., to require that she be able to tell on reflection that her belief is currently causally sustained by its grounds, together with an appropriate connecting belief to the effect that the latter evidentially support the former. (AC1), I suggest, should thus be replaced by:

(AC1)' S's belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if S can know on reflection that her belief that p is causally sustained by her beliefs B<sub>1</sub> through B<sub>n</sub>, together with the belief that these n beliefs evidentially support her belief that p.<sup>11</sup>

Now, if this much is conceded, that (AC1) must give way to (AC1)', then the first-order internalist is forced to accept at least a minimal version of second-order internalism, since (AC1)' requires *some* access to *how* the beliefs B<sub>1</sub> through B<sub>n</sub> confer justification on the belief that p. But things in fact are a bit worse for the first-order internalist, for if (AC1)' is accepted, a fuller version of second-order internalism seems naturally to follow. After all, if the subject appeals to certain grounds as evidence for her

---

<sup>11</sup> I assume that a subject's satisfying (AC1)' implies that she satisfies (AC1), insofar as her being able to tell that her belief that p is sustained by B<sub>1</sub> through B<sub>n</sub> implies that she can tell that she holds the beliefs B<sub>1</sub> through B<sub>n</sub>.

belief. then it seems she must know that the grounds do indeed evidentially support that belief. If she merely believes, without knowing, that her grounds do evidentially support her belief. then it seems entirely inappropriate for her to appeal to those grounds in support of that belief. She would appear in such a case to be in a very insecure position with respect to the truth, since she has in fact no real *reason* to think her belief true, and holds it, in effect, completely on faith.<sup>12</sup> So, (AC1)' must give way to:

(AC2) S's belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if:

- (i) S can know on reflection that her belief that p is evidentially supported by her beliefs B<sub>1</sub> through B<sub>n</sub>, and,
- (ii) S can know on reflection that her belief that p is causally sustained by her beliefs B<sub>1</sub> through B<sub>n</sub>, together with her belief that these n beliefs evidentially support her belief that p.

I shall try to argue in a moment that this new Access Condition is an acceptable replacement for the original (AC). Before turning to this, however, I should consider what the first-order internalists might say by way of response to the foregoing.

Audi has tried to argue that first-order internalism can be coherently held independently of second-order internalism by drawing a distinction between giving a justification and showing that one is justified, the suggestion being that, first, having a justified belief requires that one be in a position to give a justification for it (not to show that it is justified), and second, that one can be in a position to give a justification for one's belief by merely satisfying a first-order reflective access requirement. "*Giving a justification*," he says, "is roughly offering a ground that justifies; it is appropriate to a request for one's justification. *Showing one is justified* is roughly producing an

---

<sup>12</sup>One could ask why the subject should be required to *know* that her evidence genuinely supports her belief, as opposed to merely justifiedly believing that it does. But we have already considered this question with respect to the original formulation of the Access Condition. Recall that in Chapter 2 (pp. 51-52) I argued that a subject should be required to know that her belief is correctly inferentially based on its evidence, since she would have to be justified in believing this to the degree required for knowledge, and it would have to be true, if her belief is to be inferentially justified. A similar argument applies in the present the case.

argument with the epistemic conclusion that the belief in question is justified: it is appropriate to a challenge *of one's* justification."<sup>13</sup> Thus, he seems to think that "offering a ground that justifies" a belief requires merely that one have reflective access to the ground, and not to how it justifies, the latter being required only in order to show that one's belief is justified. Now, Audi is correct in thinking that showing that one's belief is justified does require being to some extent aware of how the grounds of that belief confer justification on it. But he is wrong if he claims, as he seems implicitly to claim here, that second-order internalism is committed to the view that having a justified belief requires that one be in a position to show that one has a justified belief, or if he claims that giving a justification requires only that one satisfy a first-order reflective access requirement. Let me briefly argue for these points.

First, while a view according to which a belief is inferentially justified only if the subject can show *that* her belief is inferentially justified would certainly be a form of second-order internalism, since in producing, as Audi says, an argument with the epistemic conclusion that her belief is inferentially justified, the subject would have to have access to all of the factors involved in the conferring of justification on the belief in question, not all forms of second-order internalism are committed to requiring this. Specifically, neither the original (AC), nor (AC2), requires a subject to tell on reflection that her belief is inferentially justified. For it does not follow from the fact that a belief is causally sustained and evidentially supported by certain other beliefs that it is inferentially justified by them, nor does this follow from the fact that this belief is correctly inferentially based on the evidence in question. Furthermore, a version of second-order internalism which does require a subject to tell on reflection that her beliefs is inferentially justified would be significantly less plausible than these more moderate conditions (or at least less plausible than (AC2)). For, we have already seen that, on pain of a vicious

---

<sup>13</sup>Audi (1993), p. 338.

regress, one cannot require a subject to have access on reflection to *all* of the factors which go into making a belief inferentially justified (she would have to tell on reflection that she can tell on reflection that these factors obtain, and so on). But if a subject is to produce on reflection an argument with the epistemic conclusion that her belief is inferentially justified, she would have to have access to all of the factors relevant to determining this.

Second. Suppose it is true, as Audi seems to suggest, that having a justified belief requires that one be in a position, not to show that it is justified, but merely to *give* a justification for it. Nevertheless, it seems quite clear that one cannot be in a position to give a justification for a belief, in the sense of citing a ground in support of it which in fact justifies that belief, unless one is aware of the ground as one's evidence, and is aware that the ground provides evidential support for the belief in question. Thus, being in a position to give a justification for one's belief requires that one not only satisfy first-order reflective access requirements, but that one also satisfy second-order reflective access requirements, contrary to Audi's suggestion.

Using the intuitive notion of being in a secure position with respect to the truth as a guide, we have built up from the first-order internalist's (AC1) a stronger condition (AC2) which requires a subject to be able to tell on reflection that her belief is evidentially supported and causally sustained by her evidentiary beliefs (together with an appropriate connecting belief). Apart from some needed modifications to be developed in the next section, I believe that this new formulation is an acceptable surrogate for the original (AC). It does, however, appear to suffer from certain difficulties which I must try to resolve.

To begin with, many theorists will doubt that the justificatory factors I have identified as reflectively accessible are genuinely capable of being known in the direct and noninferential way which we have seen is required. Consider first the relation of causal

sustenance between the belief that  $p$  and the beliefs  $B_1$  through  $B_n$  mentioned in the second condition of (AC2).<sup>14</sup> It seems to me that one does at times have a direct introspective grasp of such causal relations among one's beliefs. For example, I seem to know directly, just by focusing my attention, that my belief that the Mariners won a spring training game yesterday is sustained by my belief that I read this in the paper today, or that my belief that *The Sleepwalkers* is on the desk in the next room is sustained by my beliefs that it was there a short time ago, and that nothing could have occurred in the meantime to have caused it to move. Nevertheless, some will argue that causal relations of this kind are not manifested in consciousness in a way which allows them to be directly introspectively grasped. Just as one cannot "see" causal relations among events in the external world, but must infer their existence from constant conjunctions of observable events, so too one cannot introspect causal relations among one's own internal states. This objection raises difficult issues which I will have to postpone for the time being. I shall argue in the next chapter that, understood in the proper way, we do plausibly "see" causal relations in the external world, and that, similarly, we introspect causal relations among our beliefs. For now, I can only mention the difficulty and move on.

Consider, secondly, the relation of evidential support between the belief that  $p$  and the beliefs  $B_1$  through  $B_n$  mentioned in the first condition of (AC2). I tried to illustrate in my discussion of evidential support in Chapter 1 that we do often have direct reflective access to such relations. Thus, I considered, among others, the examples of believing that a certain candidate  $X$  will be elected because she leads in all of the major polls, and of believing that all crows are black since every observed crow has been black. It seems clear in such cases that we *know* that the premises evidentially support their conclusions,

---

<sup>14</sup>I will for convenience often omit reference to the connecting belief mentioned in this condition.

and that we know this merely by focusing our attention on the question. Some would argue, however, that, appearances to the contrary, we could not have direct reflective knowledge of evidential support relations. For, the argument proceeds, we could have reflective knowledge of evidential support relations only if such relations were, like entailment, logical connections which necessarily hold (or fail to hold) among beliefs in virtue of their contents. Since it is not plausible to think that evidential support relations are logical connections, the argument concludes, it is equally implausible to think that we have direct reflective access to them. This too is a problem I must set aside until the next chapter. I will not try there to defend the view that evidential support relations are logical connections. Rather, I shall argue that we may well have direct reflective access to evidential support relations even if such relations do not depend solely on the contents of the beliefs they relate, but depend also on certain very general nomological regularities or patterns objectively obtaining in the world.

In addition to these difficulties which I will consider in the next chapter, some theorists will worry that a condition as weak as (AC2), which requires only that a subject have reflective access to relations of evidential support and causal sustenance between her belief and its evidence, cannot adequately capture what it means to be in a secure position with respect to the truth. But I think that (AC2) still captures this intuitive notion (or at least central aspects of it) for two reasons. First, once a belief has been inferred from its evidence, its continued status as inferentially justified depends on its being causally sustained (and, of course, evidentially supported) by that evidence, together with an appropriate connecting belief. Thus we might say that in having access to the fact that her belief is causally sustained and evidentially supported by certain other beliefs, the subject has access to the factors which are *currently* responsible for her belief's being inferentially justified, and this seems enough to put her in touch with her epistemic situation in the way that being in a secure position requires.

Second, in satisfying the conditions of (AC2), a subject should be able to meet the various social expectations which explain why we desire to be in a secure position with respect to the truth. Thus, to return to the examples used to illustrate the secure position notion in the first chapter, it seems clear that if the auto mechanic believes that the car's performance falloff is due to its engine being improperly timed, and she can tell on reflection that this belief is causally sustained and evidentially supported by, say, her belief that the timing gun indicated that the #1 spark plug to be firing 10 degrees before the cylinder reaches top dead center, then she will presumably be in a position to explain to the car's owner what the problem was. For if asked to explain this, she would reflectively search her memory, and come thereby to access her belief that the timing gun indicated the #1 spark plug to be firing 10 degrees before top dead center, which (because she knows that this both provides good evidence for, and causally sustains, her belief that the car's engine was improperly timed) she would most likely proceed to cite as the cause of the car's performance falloff. Similarly, if Smith can tell on reflection that her belief that a certain individual is the best candidate for college president is sustained and supported by certain beliefs concerning this individual's qualifications, then she will be able to explain to her colleagues on the hiring committee why, in her opinion, the individual in question should become the new president of the college. For she would, in response to their queries, recall her evidentiary beliefs, and (because she can tell that they provide good evidence for her belief that so and so should be president of the college) would presumably not hesitate to use them in support of her position.

But despite the fact that (AC2) requires access to the factors which are currently responsible for a given belief's being inferentially justified, and the fact that having a capacity to access these factors will put one in a position to respond to challenges and requests for evidence, some theorists might suggest that stronger conditions would more adequately capture what it means to be in a secure position with respect to the truth.

First of all, it might be argued, if a subject's having access to the fact that her belief is evidentially supported and causally sustained by certain others is to put her in a position to respond to challenges to that belief (i.e., to be in a secure position), then she must appreciate the epistemic significance of these factors, that they help to make her belief reasonable. While it cannot be required that the subject know on reflection that her belief is inferentially justified, she must at least appreciate that the justificatory factors in question, the relations of evidential support and causal sustenance between her belief and its evidence, *contribute* to her belief's being inferentially justified. Thus, it might be suggested that (AC2) be replaced with:

(AC3) S's belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if:

- (i) S can know on reflection that her belief that p is evidentially supported by her beliefs B<sub>1</sub> through B<sub>n</sub>, and,
- (ii) S can know on reflection that her belief that p is causally sustained by her beliefs B<sub>1</sub> through B<sub>n</sub>, together with her belief that these n beliefs evidentially support her belief that p, and,
- (iii) S can know on reflection that her belief that p's being evidentially supported and causally sustained by her beliefs B<sub>1</sub> through B<sub>n</sub> (together with her belief that these n beliefs evidentially support her belief that p) contributes to her belief that p being inferentially justified.

However, the suggested condition (iii) is ambiguous, and on either of its readings it appears to be unacceptable. In saying that the subject must be able to tell on reflection that the factors mentioned in conditions (i) and (ii) of (AC3) contribute to her belief's being inferentially justified, this could mean that she is aware that these factors are specified in some philosophical account of inferentially justified belief, or just that they help her belief to be reasonable or warranted in some intuitive, everyday sense. The former is implausible, since ordinary people do not operate with philosophical accounts of inferential justification (i.e., they do not operate with explicit lists of precisely formulated conditions which they understand to be "conditions of inferential justification"). But condition (iii) is unreasonable on the latter reading as well, for on that understanding it suffers from a certain form of circularity. If (AC3) is to be of any

interest to us at all, then it must be offered as an explication of the intuitive notion of being in a secure position with respect to the truth, that is, as an explanation of the everyday notion of reasonable or warranted belief. But this everyday notion appears in condition (iii) on the indicated second reading. At the very least, this limits its usefulness as an explication of this intuitive notion.

There is, however, a much more plausible alternative to (AC3) which is still stronger than (AC2). Recall the argument given above to the effect that a subject's belief is inferentially justified by her evidentiary beliefs only if she has access to the fact *that* these beliefs serve as her evidence. It might be suggested that the way in which I cashed this requirement out, specifically, that the subject must have access to the fact that her belief is causally sustained by her evidentiary beliefs (together with an appropriate connecting belief), does not adequately capture what it intuitively means for the subject to have access to the fact that certain beliefs serve as her evidence. Knowing that certain beliefs serve as her evidence requires that she also *remember*, in a way which yields direct knowledge, having inferred her belief *from* these other beliefs. To illustrate, consider again the case of the auto mechanic. If she is to truly know that her belief that the timing gun showed the #1 spark plug to be firing 10 degrees before top dead center is what serves as her evidence for thinking that the car's engine is improperly timed, then she must not only know that the latter belief is causally sustained and evidentially supported by the former, but that she did indeed at one time *infer* the latter from the former. It might be thought, then, that the Access Condition should require a subject to have access not only to the fact that her belief is sustained and supported by certain other beliefs, but also to the fact that she inferred the former from the latter at one time. If we assume that a "memory impression" of having inferred a belief from certain other beliefs can yield direct knowledge of that inferring event (for we have seen that only directly

knowable justificatory factors may be required to be reflectively accessible), then (AC2) might be replaced with:

(AC4) S's belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if:

- (i) S can know on reflection that her belief that p is evidentially supported by her beliefs B<sub>1</sub> through B<sub>n</sub>, and,
- (ii) S can know on reflection that her belief that p is causally sustained by her beliefs B<sub>1</sub> through B<sub>n</sub>, together with her belief that these n beliefs evidentially support her belief that p, and
- (iii) S has direct memory knowledge of having inferred her belief that p from her beliefs B<sub>1</sub> through B<sub>n</sub> at some earlier time.

I admit that (AC4) has considerable attraction relative to the aim of explicating what it means to be in a secure position with respect to the truth, since it requires access to a further factor which is crucially involved in a belief's being correctly inferentially based on its evidence. I suspect, however, that satisfying condition (iii) of (AC4) is not necessary for being in a secure position with respect to the truth. It strikes me as superfluous, for example, to require the auto mechanic to have direct memory knowledge of having inferred her belief that the car's engine is improperly timed from her evidentiary belief, since she will be able to explain the problem to the car's owner even if some time has passed and she has failed to retain a specific memory impression of the inferring event in question. While she would, in such a circumstance, presumably still "remember" that she inferred her belief that the car's engine is improperly timed from her belief that the timing gun indicated the #1 spark plug to be firing 10 degrees before top dead center, this would not constitute direct memory knowledge of the sort required by condition (iii) of (AC4) if she no longer recalls the specific incident in which this inferring took place.

So, despite the attractiveness of (AC4), I propose to put it aside, at least provisionally, and continue to operate with (AC2) as the best reformulation of the original (AC). As we have seen, (AC2) preserves some measure of a requirement that a subject have access to *how* her evidence confers justification, and the factors required to be thus

accessible appear, pending the fuller discussion of the issue in Chapter 5, to be capable of being noninferentially known to obtain.

#### 4. Refining the Access Condition.

At the beginning of this chapter, I identified two questions which I would try to answer. First, what justificatory factors should be required to be reflectively accessible in the proper formulation of the Access Condition? And second, what does it mean to say that a subject *can* access these factors on reflection? I defended an answer to the first question in the previous section. Specifically, I argued for the following reformulation of (AC):

(AC2) S's belief that p is inferentially justified by her beliefs B<sub>1</sub> through B<sub>n</sub> only if:

- (i) S can know on reflection that her belief that p is evidentially supported by her beliefs B<sub>1</sub> through B<sub>n</sub>, and,
- (ii) S can know on reflection that her belief that p is causally sustained by her beliefs B<sub>1</sub> through B<sub>n</sub>, together with her belief that these n beliefs evidentially support her belief that p.

I must now consider the second question.

On the weakest interpretation of 'can,' (AC2) requires only that it be *logically possible* for the subject to tell on reflection that her belief that p is causally sustained and evidentially supported by certain other beliefs. But being in a secure position with respect to the truth would seem to require more than this. It is not enough that there be *some* possible circumstance in which the subject is able to access the factors specified in (AC2); there must be at least a substantial likelihood of her becoming aware of the fact that her belief is sustained and supported by its evidence. At the same time, it is clear that we must acknowledge the possible existence of circumstances or factors which would prevent a subject from having the required access to the relevant justifying factors, without necessarily implying that, in such circumstances, the subject's belief would not be inferentially justified. Thus, given adequate time to reflect, and kept free of

distractions, a subject might be quite adept at accessing her evidence, but if she is hungry, or tired, or in a hurry, she may not at that time have the required reflective access to her evidence. (AC2), therefore, must be modified to include an “appropriate circumstances” clause<sup>15</sup> which allows that, in at least some situations, the subject may not be able to tell on reflection that her belief that *p* is supported and sustained by certain evidentiary beliefs. The problem is to incorporate this clause in a proper way. For example, the following would not be an acceptable way to incorporate the appropriate circumstances clause:

(AC5) *S*'s belief that *p* is inferentially justified, at a time *t*, by her beliefs *B*<sub>1</sub> through *B*<sub>*n*</sub> only if:

- (i) provided that circumstances are appropriate at *t*, *S* can at *t* come to know by reflection that her belief that *p* is evidentially supported by her beliefs *B*<sub>1</sub> through *B*<sub>*n*</sub>, and,
- (ii) provided that circumstances are appropriate at *t*, *S* can at *t* come to know that her belief that *p* is causally sustained, at *t*, by her beliefs *B*<sub>1</sub> through *B*<sub>*n*</sub>, together with her belief that these *n* beliefs evidentially support her belief that *p*.<sup>16</sup>

To see that this version of the Access Condition is not acceptable, imagine what consequence it would have if circumstances were not, at *t*, appropriate for reflection. If circumstances are not appropriate for reflection at a given time *t*, then, by (AC5) the subject is, in effect, exempt from satisfying the Access Condition at *t*, for (i) and (ii) only require reflective access to their respective justificatory factors *if* circumstances are indeed appropriate at *t*. But this would allow a subject's belief to be inferentially justified at *t*, even if she is not at all likely to satisfy the Access Condition *when* circumstances are appropriate. It is much more reasonable to suppose that inferential justification requires that if circumstances are *not* appropriate for reflection, then if they *were* appropriate, the subject would be able to tell on reflection that the relevant justificatory factors or

---

<sup>15</sup>I will not try to give a precise specification of the range of appropriate circumstances, but will rely on the reader's intuitive understanding of this notion.

<sup>16</sup>As we are nearing a final formulation of the Access Condition, I will begin to mention required temporal indices explicitly.

conditions obtain. If we allow that a proposition of the form “Were it the case that *p*, then it would be the case that *q*” need *not* imply that it is not the case that *p*, then we might try to modify (AC5) so as to include an appropriate circumstances clause in the following way:

- (AC6) *S*'s belief that *p* is inferentially justified, at a time *t*, by her beliefs *B*<sub>1</sub> through *B*<sub>*n*</sub> only if:
- (i) were circumstances appropriate at *t*, *S* would if she reflected on the matter at *t*, come to know that her belief that *p* is evidentially supported by her belief *B*<sub>1</sub> through *B*<sub>*n*</sub>, and,
  - (ii) were circumstances appropriate at *t*, *S* would if she reflected on the matter at *t*, come to know that her belief that *p* is causally sustained at *t* by her beliefs *B*<sub>1</sub> through *B*<sub>*n*</sub>, together with her belief that these *n* beliefs evidentially support her belief that *p*.

This version of the Access Condition, while closer to being correct than previous versions, nevertheless stands in need of a certain improvement. The problem concerns the requirement that *S would* in appropriate circumstances tell on reflection that her belief that *p* is sustained and supported by her beliefs *B*<sub>1</sub> through *B*<sub>*n*</sub>. It seems at least conceivable that a subject might have an inferentially justified belief and yet *fail* in an appropriate circumstance to access the relevant justificatory factors. This possibility is convincingly suggested in the following passage by Bonjour:

It is clear that we often think, commonsensically, that we know various things which would have to be justified on an inferential basis if justified at all, but for which the required inferences not only have not been formulated explicitly, but could not be formulated without very considerable reflective effort--and perhaps not even then.

For example, I believe that the piece of paper on which I am now typing is the very same piece of paper on which I was typing late yesterday afternoon. I am confident that this belief is justified; and it seems inescapable that if it is indeed justified, its justification is inferential in character. But I am not at all sure, at first glance at least, how the justificatory inference would go. Some of the premises are obvious enough, but any adequate argument would have to appeal to general principles concerning the individuation of physical objects which I simply do not know at present how to state. If I were to spend some time reflecting on the matter, I would no doubt be able to come closer to an adequate formulation of the argument and *might* even succeed more or less completely. But it also seems clear that any actual attempt might fall seriously

short of complete success.<sup>17</sup>

The thrust of Bonjour's argument here is that there are cases in which the evidence or reasoning which supports a belief is so complicated or subtle that the subject *might* fail recall it, even after "considerable reflective effort."<sup>18</sup> We might question whether the case Bonjour describes really fits this description. The source of Bonjour's doubts concerning his ability to reflectively formulate the reasons which support his belief is that such reasons would have to involve, he thinks, principles for individuating physical objects which he is presently unable to state in precise terms. But to what extent must he be able to precisely formulate such principles in order to be in a secure position with respect to the truth? It seems enough for him just to say that the piece of paper looks exactly like the one on which he was typing the day before (it bears all of the words which he typed then, and none which he did not, is situated in the typewriter exactly as he had left the paper the day before, and so on), and that it is highly improbable that any person, or any fortuitous process, could have caused some *other* piece of paper (bearing the same words) to have taken the place of the previous day's paper. Perhaps Bonjour is thinking, rightly, that this would not be enough to satisfy a skeptic concerning our knowledge of the external world. But, as we discussed in the first chapter, being in a secure position with respect to the truth does not require that one be able to answer skepticism.<sup>19</sup>

Still, it does seem that normal people might not always be able to recall evidence for their beliefs, even in appropriate circumstance, as Bonjour claims. Imagine a defense attorney whose reasons for thinking that his client is innocent are so convoluted that he sometimes has trouble formulating them in a perspicuous way, and occasionally fails to recall certain crucial elements in the argument. He might in presenting his defense to the

---

<sup>17</sup>BonJour (1985), p. 20.

<sup>18</sup>Ibid., pp. 19-20.

<sup>19</sup>This point will be treated further in the next chapter.

jury suffer from such a memory lapse, or be unable to adequately express his reasoning, not because he is nervous or distracted or because circumstances are in some other way inappropriate, but simply because he *happens* to inadequately recall his argument. While this failure on his part would likely have detrimental consequences for his client, and no doubt would affect some people's judgment of his abilities as an attorney, it would not seem to affect the *epistemic* status of his belief in his client's innocence. For he does *have* reasons which he could adequately recall in many circumstances. He just does not succeed in doing so on this occasion.

So, it seems to strong to a requirement to hold, as (AC6) does, that in appropriate circumstances a subject *would* come to know on reflection that her belief is sustained and supported by her evidence. Bonjour suggests as an alternative that a subject need only have what he calls *tacit awareness* of justificatory factors, so that she could *in principle* access them. If we think, however, of paradigm cases of subjects who are in secure positions with respect to the truth, it seems that such individuals have more than a tacit awareness of the relevant justificatory factors. The auto mechanic is presumably able to recall the reasons for her belief that the engine is improperly timed without undue effort. Similarly, the grounds of Smith's belief that such and such an individual ought to be the new president of the college would appear to be fairly readily available to her. It seems to me, then, that while being in a secure position with respect to the truth may not require that a subject would in any appropriate circumstance be able to access the relevant justificatory factors, it does require that there be at least a *high probability* of the subject's accessing those factors, though, of course, depending on the complexity of her reasons, some reflective effort on the subject's part might be involved. This suggests a modification to (AC6) which, I believe, yields the correct formulation of the Access Condition:

(AC\*) S's belief that p is inferentially justified, at a time t, by her beliefs B<sub>1</sub> through B<sub>n</sub> only if:

- (i) were circumstances appropriate at t, it is very probable that, with perhaps some effort, S would, if she reflected on the matter at t, come to know that her belief that p is evidentially supported by her beliefs B<sub>1</sub> through B<sub>n</sub>, and,
- (ii) were circumstances appropriate at t, it is very probable that, with perhaps some effort, S would, if she reflected on the matter at t, come to know that her belief that p is causally sustained at t by her beliefs B<sub>1</sub> through B<sub>n</sub>, together with her belief that these n beliefs evidentially support her belief that p.

The Access Condition in this final version which I shall accept is still subject to certain questions and objections. The most pressing of these is the issue of whether the justificatory factors specified in conditions (i) and (ii), i.e., the relations of sustenance and support between a belief and its evidence, are the sorts of things which can be directly or noninferentially known to obtain. I shall consider this in the next chapter.

In closing, I should comment on a potentially puzzling feature which emerges when this new formulation of the Access Condition is combined with the Basis Condition as stated at the end of the previous chapter. One aspect of the latter condition, as we have seen, is that the subject is required to hold a connecting belief to the effect that her inferred belief is evidentially supported by the beliefs from which it was inferred. Thus, correct inferential basing appears to require a subject to be *in some sense* already aware that her inferred belief is evidentially supported by the beliefs from which it was inferred, independent of satisfying the Access Condition. Thus, it may not be immediately clear to the reader what clause (i) of (AC\*), which requires that the subject know on reflection that her inferred belief is evidentially supported by the beliefs from which it was inferred, adds to the connecting belief requirement in (BC\*). But, in fact, the appearance of redundancy here is illusory. For the connecting belief which the subject is required to hold in satisfying the Basis Condition need *not* be epistemically justified or known to be true. All that is required by the Basis Condition is that the subject *believes*

that her belief that  $p$  is evidentially supported by her beliefs  $B_1$  through  $B_n$  and that this connecting belief play an appropriate role in guiding the inference which produces her belief that  $p$ . The subject need not, in satisfying the Basis Condition, have reflected on whether her belief that  $p$  is indeed evidentially supported by her beliefs  $B_1$  through  $B_n$ , and thus she may not in fact be justified in believing that  $p$  is so supported, and in consequence may not *know* that it is evidentially supported. In order, however, to satisfy the *Access Condition*, and not merely the Basis Condition, the subject must be capable of coming to know on reflection, that is, by focusing her attention on the question, whether her belief that  $p$  is evidentially supported by her beliefs  $B_1$  through  $B_n$ . Putting this another way, if a subject is to jointly satisfy (BC\*) and (AC\*), then she must in a sense have two modes of access to the evidential support relation between her belief that  $p$  and her beliefs  $B_1$  through  $B_n$ . For, she must, first, believe that the latter evidentially support the former, and second she must be able to come to *know* this fact (and not simply believe it) on reflection.

#### 4. Summary.

In this chapter I developed a reformulation of the Access Condition which, together with the reformulated Basis Condition offered in the third chapter, goes a long way toward providing a concrete and detailed account of inferential justification. The reformulated Access Condition proves to be significantly weaker than might have been expected from its initial formulation. In its final form, the Access Condition requires, not that the subject tell on reflection that her belief is correctly inferentially based on others, but only that certain conditions obtain which are necessary for it to be correctly inferentially based on other beliefs. Specifically, the Access Condition requires, in the end, merely that, in appropriate circumstances, it is very probable that the subject will be able to know on reflection that her belief is evidentially supported by certain other beliefs,

and that it is causally sustained by them (together with an appropriate connecting belief).

I argued that this weakened formulation still captures the intuitive notion of being in a secure position with respect to the truth, since the justificatory factors which it requires to be reflectively accessible figure most prominently in the explanation of why an inferentially justified belief continues to be justified after it is formed, and since a subject's having a capacity to access these factors will enable her to meet everyday challenges to her beliefs.

## Five

### Responses to Objections

#### *1. The Theory Reviewed.*

In the previous chapters, I have developed part of an internalist theory of inferential justification based on what I have termed a *secure position* rationale. The theory I have offered rests on the idea that a belief's being inferentially justified is a matter of the subject being in a secure position with respect to the truth of that belief, meaning roughly that the subject holds the belief on the basis of adequate evidence, and has that evidence, as it were, "at her fingertips"; she can access or recall that evidence as the situation demands. The account offers a partial explication of this intuitive desideratum in terms of two necessary conditions. The first of these, called the Basis Condition, requires beliefs to be held on the correct inferential basis of evidentiary beliefs, and specifies several factors necessary for this basing relationship to obtain. The second, called the Access Condition, requires that, in the right circumstances and with high probability, two of the factors necessary for correct inferential basing be accessible to the subject on reflection. These conditions, expressed together in their final formulations, yield the following (partial) theory of inferential justification (IJ):

- (IJ) S's belief that *p* is inferentially justified, at a time *t*, by her beliefs  $B_1$  through  $B_n$  only if the following conditions are satisfied:

*(The Basis Condition)*

- (i) The sequence of beliefs  $\langle B_1, \dots, B_n, p \rangle$  is of a kind *K* such that the first *n* members evidentially support the final member.
- (ii) S believes that her beliefs  $B_1$  through  $B_n$  evidentially support her belief that *p*.
- (iii) S has the following disposition *D*: If S has the first *n* beliefs in sequence of kind *K*, and S believes that these *n* beliefs evidentially support some further belief which would complete a sequence of kind *K*, and these

beliefs are simultaneously activated, then S is disposed to form the further belief which would complete a sequence of kind K.

- (iv) There was a time  $t'$ , not later than  $t$ , such that S held at  $t'$  the beliefs  $B_1$  through  $B_n$ , and believed that these  $n$  beliefs evidentially support the belief that  $p$ , and S came to believe that  $p$  at  $t'$  as a result of her disposition  $D$ , triggered by simultaneous activation of her beliefs  $B_1$  through  $B_n$  together with her belief that these  $n$  beliefs evidentially support the belief that  $p$ .
- (v) For any time  $t''$ , no earlier than  $t'$  and no later than  $t$ , S's belief that  $p$  is sustained at  $t''$  by her beliefs  $B_1$  through  $B_n$ , together with her belief that these  $n$  beliefs evidentially support her belief that  $p$ .

*(The Access Condition)*

- (vi) Were circumstances appropriate at  $t$ , it is very probable that, with perhaps some effort, S would, if she reflected on the matter at  $t$ , come to know that her belief that  $p$  is evidentially supported by her beliefs  $B_1$  through  $B_n$ .
- (vii) Were circumstances appropriate at  $t$ , it is very probable that, with perhaps some effort, S would, if she reflected on the matter at  $t$ , come to know that her belief that  $p$  is causally sustained at  $t$  by her beliefs  $B_1$  through  $B_n$ , together with her belief that these  $n$  beliefs evidentially support her belief that  $p$ .

As I have said, this account does not completely explain or capture the intuitive idea of a secure position for a number reasons. First, no analysis has been offered of the notion of evidential support relations which figures so prominently in the theory. Some suggestion has been made that evidential support is objectively truth-linked, in that there presumably is in normal circumstances a close relationship between a belief's being evidentially supported and its being objectively very probably true. But it is unclear exactly what *sort* of objective probability is involved in a belief's being evidentially supported (whether, that is, the probability is logical or nomological in character), and it is equally unclear what the exact relationship *is* between evidential support and objective probability (i.e., it is unclear whether these quantities are in all cases identical).

Second, no account has been offered of the nature of reflective access. All I have claimed (and all I will continue to claim, with any confidence) is that having reflective access to a state of affairs involves being in a position to know on reflection, that is, by carefully focusing one's attention on the matter, that the state of affairs in question does

indeed obtain. But a complete understanding of the factors involved in being in a secure position with respect to the truth requires a fuller investigation into how and why the process of careful reflection, in the proper circumstances, yields knowledge.

Finally, the account as presented just now, in the form (IJ), gives, as I have said all along, only necessary conditions of a belief's being inferentially justified. As I discussed in Chapter 2, the evidentiary beliefs which confer justification on an inferentially justified belief must themselves be justified, raising questions about the overall justificatory structure of the subject's belief system, whether the evidentiary beliefs must be justified independently of the target belief, thus yielding a foundational structure, or whether the subject's beliefs can be justified ultimately by their being embedded in a coherent web of mutually supportive inferential relationships. In addition to this question about the justification of evidentiary beliefs, there is the further issue, also identified in Chapter 2, of the role of potential defeating evidence, that is, beliefs or evidence which the subject failed to take into account in forming a belief, but which if she had would have led her to withhold that belief. One might say, then, that (IJ) focuses only on *local* factors in inferential justification, in the sense that it describes only the immediate connections between a belief and its evidence which operate in conferring justification on the former. Global factors of justification, i.e., further features of the subject's larger belief system relevant to determining justification, have not been addressed in detail.

Despite the incompleteness of (IJ) as an account of what it is to be in a secure position with respect to the truth, it does succeed in capturing important aspects of this pretheoretic cognitive desideratum, as I have tried to argue. In the first place, it is surely part of being in a secure position with respect to the truth that one's beliefs be, to the fullest extent possible, correctly inferentially based on adequate evidence. It is not enough that one simply "have" evidence; one must actually *use* evidence, take it into account, in

forming beliefs. And once evidence has been employed in forming a given belief, one is surely in a more secure position if one retains that evidence, and continues to hold the belief on the basis of it, than if one allows the evidence to fade from memory.

In the second place, being in a secure position in holding a particular belief ought to require that one have reflective access to the most salient relations by which the relevant evidentiary beliefs confer justification on that belief. I have argued that with respect to the inferential justification of a given belief at a particular time, the most salient relations operative in the conferring of justification are the relations of evidential support and causal sustenance between the belief and its evidence, i.e., the relations which are required to be reflectively accessible by the Access Condition component of (IJ).

So, I think that, as far as it goes, (IJ) adequately explicates the intuitive notion of being in a secure position with respect to the truth, and is, to that extent, a *correct* partial theory of inferential justification. The task now is to defend this account against potential objections. I will consider here five objections, similar versions of which have been directed at previous internalist theories. I divide these objections into two categories. Those in the first allege that the account offered in (IJ) is not satisfiable *in principle*, those in the second that it is not widely satisfied in practice. Since a great many of our beliefs about the past, future, and things unobserved must be inferentially justified if justified at all, the thrust of these objections is that ordinary adult human beings have very few (or no) justified beliefs on the present account of inferential justification, a result which seems to count strongly against the correctness of (IJ).

The Access Condition is the target of the objections in the first category (since there appears to be no reason in principle why the Basis Condition could not be satisfied in a given case). It is argued that neither relations of causal sustenance nor relations of evidential support are the *sort* of things which could be reflectively accessed. I will consider this question first in regard to relations of causal sustenance, where the problem

(and the required solution) can be seen most clearly. The difficulty in regard to introspective access to causal connections is exactly analogous to the difficulty Hume discussed in regard to observational access to causal connections, namely, that there is never any discriminable element of the phenomenal field corresponding to a causal relation. One apparently observes or introspects only the causally connected objects or events (or states, in the case of beliefs), without seeing or introspecting the causal connection between them—the existence of *that* must be inferred from what one *does* see or introspect. As a response, I claim that on a “naturalistic” conception of observational knowledge as belief produced by reliably or properly functioning sensory mechanisms, one does on many occasions have observational access to causal connections between objects or events, e.g., between the motions of a cue ball and three ball on a pool table, despite the fact that causal connections are not phenomenally manifested in the visual field in the same way as the shapes and colors of billiard balls. This suggests that it is only on a very particular and mistaken conception of introspection that one’s sphere of introspective access is limited to those items which are directly phenomenally manifested in consciousness. I argue that since this “Cartesian” conception of introspection, as I call it, cannot account for our introspective knowledge of our own belief states, it cannot be correct, thus paving the way for a naturalistic account which makes sense both of our introspective access to beliefs *and* to the causal relations between them.

I consider next the objection that reflective noninferential knowledge of objectively truth-linked relations of evidential support is impossible. Here the argument is that such relations would be reflectively accessible only if *logical* in character, that is, entirely and essentially a function of the *contents* of the beliefs or propositions thus related, so that they could not *be* those very beliefs or propositions were they not evidentially related in that way. But, it is claimed, the notion of purely logical objectively truth-linked evidential support relations is extremely dubious. If a set of beliefs or

propositions evidentially supports a further belief or proposition in way that renders the latter very probably true, given the truth of the former, this would seem to depend not merely on the contents of the propositions related, but on certain very general nomological regularities or patterns objectively obtaining in the world which seem not to be reflectively accessible. Once again, I argue in response that a naturalistic account of our knowledge of evidential support relations (as opposed to a “Cartesian” account in terms of the direct cognitive grasp of necessary connections) may be able to make sense of such knowledge, even if those relations are nomological rather than logical in character.

The key, then, to answering the first two objections is to recognize that they each presuppose a particular Cartesian account of the relevant reflective faculty, or at least presuppose that internalist views are committed to such a Cartesian account. This brings us, therefore, to the third objection, to wit, that internalist theories cannot avail themselves of the naturalistic account of reflective knowledge required to make sense of our access to evidential and causal relations among beliefs. I shall argue that the secure position rationale of the present theory is not lost if the shift is made to a naturalistic account of reflective knowledge.

The last two objections I will consider belong in the second category identified above, i.e., they maintain not that the account is unsatisfiable in principle, but that it is not widely satisfied in practice. These objections turn on a wealth of empirical data concerning ordinary human deficiencies in information storage capacity and reflective power which seem to indicate that people routinely forget the evidence from which their beliefs are derived, or are unable to recall or access that evidence except in highly artificial, and thus extremely rare, circumstances. The response I will try to develop is that the Basis and Access Conditions represent slightly idealized standards of justification which are nevertheless useful for epistemic evaluation in those situations in which

individuals implicitly strive for a high level of epistemic achievement--in other words, those situations in which they strive to be in a secure position with respect to the truth.

I will now consider the first category of objections in more detail.

## 2. *Is The Account Satisfiable in Principle?*

In this section, I will consider the alleged impossibility of reflectively accessing relations of causal sustenance and evidential support. I shall argue that on a naturalistic account of reflective knowledge these relations are accessible in principle, and that such a naturalistic account is not incompatible with my theory of inferential justification.

*Objection 1.* Causal relations among beliefs, in particular relations of current causal sustenance, are not reflectively accessible.

Robert Audi offers a clear statement of the reasoning which underlies this objection:

Consider an (apparently) justified belief that one will go to the office tomorrow. Can one, by introspective reflection, become aware of the relevant kind of causal sustaining relation between one's grounds for this and one's believing it?

A negative answer finds support along these lines. First, there seems to be no standard set of ways that such causal connections manifest themselves in consciousness: one introspects at best a constant conjunction. Moreover, it appears that to know or justifiably believe that my belief that I will go to the office is causally sustained by, say, my belief that nothing will prevent realization of my intention to go, I must know such things as this: that if I lacked the latter belief, I would not hold, or would less strongly hold, the former. Can one know or justifiably believe such counterfactuals through introspective reflection? This is far from obvious. Indeed, it is not even obvious that when believably entertaining the premises of a valid syllogism produces one's belief of the conclusion, one can know or justifiably believe, on the basis of introspection, that this production relation actually holds.<sup>1</sup>

---

<sup>1</sup>Audi (1993), pp. 335-336. In fairness to Audi, he himself believes that we *do* at times have introspective access to causal connections. Though he is not as explicit as he could be on the question, I believe he would be receptive to some of the suggestions I make in my response to this first objection.

These remarks are analogous to some of Hume's famous observations on the impossibility of perceiving causal connections among external objects or events. When a billiard ball strikes another, sending it in motion, one is, according to Hume, unable to perceive any connection between the two balls distinct from one's perception of the first ball coming into contact with the second, and the subsequent motion of the latter. "The first ball is in motion; touches the second; immediately the second is in motion: and when I try the experiment with the same or like balls, in the same or like circumstances, I find, that upon the motion and touch of the one ball, motion always follows in the other. In whatever shape I turn this matter, and however I examine it, I can find nothing further."<sup>2</sup> Hume perceives "nothing further," nothing other than the successive motions of billiard balls, because (as we might put it) no discriminable phenomenal element of the visual field corresponds to a causal relation between the billiard balls. The causal relation is thus not something *seen* in the way that the billiard balls themselves are seen.

The argument expressed by Audi holds for similar reasons that causal relations are not introspectively accessible. While one is able to introspectively access, say, a stabbing pain in one's elbow, or an imagined vivid blue patch (these items, as it were, are discriminable elements of the introspected phenomenal field), one is unable to introspect causal connections between internal states. Since causal relations are never discernible phenomenal elements of the introspective field, "one introspects at best a constant conjunction."

*Reply.* The objection goes to the heart of the nature of observation and introspective knowledge. A truly adequate response would require worked out theories of these sorts of knowledge, a project which is beyond the scope of this work. The most I can hope to do here is to sketch the outlines of a conception of introspective and observational knowledge which will illuminate the mistake underlying the first objection.

---

<sup>2</sup>Hume (1938), p. 12.

Implicit in Hume's remarks quoted above is a view of perceptual knowledge according to which one can know observationally that a given item has a particular property, or stands in a particular relation to some further object only if the property or relation in question is manifested in an identifiable phenomenal feature of the perceptual field. One can know observationally that the cue ball causes the three ball to fall into the corner pocket only if the causal relation is a phenomenal element alongside of, and in roughly the same way as, the phenomenal elements corresponding to the cue and three balls. Since, in fact, the causal connection does not appear in the visual field in this way, one's knowledge of it must not be observational, so the argument goes, but inferential: one infers from the successive motions of the balls (and from similar past successions of motions), that the one causes the other to move.

Now, leaving aside for the moment the question of whether causal connections are directly observable, it is clear that the conception of observational knowledge just discussed cannot be quite correct. For we do at times know that objects have certain properties, or stand in certain relations, even when the properties or relations in question are not represented by simple discriminable elements of the phenomenal field. I may know observationally that the cue ball is to the left of the three ball, but there would appear to be no discriminable phenomenal element corresponding to the relation of *being to the left of*. One might argue, perhaps, that I can see that the balls are separated by a small expanse of green felt, and the patch of greenness is the phenomenal element corresponding to the *being to the left of* relation. But this element could not serve to visually distinguish the *being to the left of* relation from an indefinite number of other spatial relations between the two balls, such as, for instance, their being within two feet of each other, or the cue ball's being to the *right* of the three ball.<sup>3</sup> To take a second

---

<sup>3</sup>We can also see that the expanse of green is to the left of the three ball. There is no expanse of green corresponding to this relation.

example, one might *see* that the two balls are in physical contact, but there is no discriminable phenomenal element corresponding to their relation of spatial contiguity. Here it is no good identifying the *absence* of an expanse of green as the relevant element of the visual field; it would seem odd to speak of a discernible phenomenal element corresponding to the absence of some further phenomenal element.

Or, consider the process of perceptually identifying an object as a cue ball or a three ball, or anything else, for that matter. I might perceptually recognize and identify an object before me as a car. My belief that the item is a car *seems* both visual and noninferential, but the item's *being a car* does not correspond to an element in my visual field in the same way as the distinctive shape and color it presents. One might argue that the fact that I am able to identify the object before me as a car depends on prior experience with items of this type, and that I am in effect *inferring* that the item is a car on the basis of previously acquired information to the effect that items which look in such and such ways are cars. But I want to advocate a different view, according to which my perceptual mechanisms are *conditioned* by past experience, so that I would be unable to perceptually recognize the item in question as a car without the requisite prior experience (and the attendant information processing) by which I learned to discriminate cars from noncars. But these past experiences (or the stock of beliefs they yielded), which make possible the formation of my belief that this is a car, do not serve as premises from which I *infer* that this is a car. Once the distinctive characteristics of cars have been learned and embedded in my conceptual system, the appropriate sensory input triggers my belief that this is a car, without mediation from other beliefs (a similar story could be told in the case of my belief that the cue ball is to the left of the three ball, or that they are spatially contiguous).

We could make sense out of a view of this latter sort, according to which perceptual beliefs are conditioned by, but not necessarily inferred from, previously

acquired information, from a naturalistic or causal perspective on perceptual knowledge, i.e., a conception of such knowledge as belief produced in the right manner, presumably by being grounded in an appropriate sensory experience<sup>4</sup> in such a way that the belief is very likely to be true. If my belief that there is a car before me is produced by reliably functioning visual mechanisms from appropriate sensory input, then, perhaps, my belief counts as knowledge.<sup>5</sup> Of course, my visual mechanism will not be reliable in this circumstance unless I have learned to distinguish cars from noncars on the basis of visual clues. But this training serves only to create the disposition to form beliefs that cars are present when I have certain sensory experiences. Once this disposition is in place, there is no reason to suppose that the production of the visual belief depends on any mediating premises concerning my past experience with cars.<sup>6</sup>

If we adopt the suggested causal or naturalistic view of perceptual knowledge, then it becomes plausible (or at least possible) to think that we have direct, noninferential perceptual knowledge of causal relations, in at least some circumstances. If I see that the cue ball has contacted the three ball, which is now in motion, I may form the belief that the cue ball has caused the three ball to move, and this may be formed reliably and directly, that is, noninferentially, from the perceptual input, without mediation from vast numbers of prior observations of the transmission of motion by contact. Alternatively, to focus on a case of causal sustenance, I may believe that the table is sustaining the book I see before me, and again, the belief may be formed in a reliable way from the current visual cues. Of course, prior experiences with sustaining relationships may be required in order for me to *recognize* that the present situation provides an instance of causal

---

<sup>4</sup>It may be that all observational knowledge must be based on or grounded in an appropriate sensory experience, even if the property or relation thereby known to obtain does not correspond in any straightforward way to a feature of that experience.

<sup>5</sup>Or, if the reliability of the relevant mechanisms is insufficient for the resulting belief to count as knowledge, then perhaps, as Plantinga (1993b) would suggest, it is the *proper functioning* of those mechanisms which is knowledge yielding.

<sup>6</sup>See Alston (1991), p.91. for a similar point.

sustenance. But this does not show that I have inferred my belief that the table sustains the book from mediating premises about these past experiences.<sup>7</sup>

Let us turn now to the question of introspection. Here we might distinguish two views analogous to the views about perception just discussed. On one view, a subject can introspect only discriminable phenomenal elements currently manifested in consciousness (blue patches, entertained images, dull aches, stabbing pains, piercing sounds, and so on). On the other view, a subject has introspective access to mental states, events, or items (and relations between them), whenever she has beliefs about those states, events, or items which are produced in the right manner, presumably by being reliably or properly grounded in appropriate introspectable phenomenal elements, in much the same way as perceptual beliefs are grounded in appropriate sensory experiences.<sup>8</sup> According to this second view, introspective knowledge is not different in kind from perceptual knowledge of the external environment. Both sorts of knowledge, perceptual and introspective, are seen to depend on the same basic factors, namely, the proper production of perceptual or introspective belief by way of appropriate nondoxastic inputs. As D. M. Armstrong has suggested,

[P]erception, considered as a mental event, is the acquiring of information or misinformation about our environment. It is not an 'acquaintance' with objects, or a search light that makes contact with them, but is simply the getting of beliefs. Exactly the same must be said of introspection. It is the getting of information or misinformation about the current state of our mind.<sup>9</sup>

The first view of introspection rejects this idea. Whatever one must say about perception, this view maintains, introspection does consist in "acquaintance with," or direct grasping

---

<sup>7</sup>It is quite possible, of course, that some causal judgments are innate and are thus not even conditioned by past experience, let alone inferred from it. For simplicity, however, I will continue to ignore this possibility.

<sup>8</sup>Again, such knowledge may be grounded in phenomenal elements even if the properties or relations thereby known to obtain do not correspond to such elements in a straightforward manner.

<sup>9</sup>Armstrong (1968), pp. 325-326.

of, currently manifested conscious phenomenal elements. Thus, it is claimed, introspective knowledge is different in kind from perceptual knowledge, and depends upon entirely different factors. I will call this view, that introspection is limited to the direct grasping of currently manifested discriminable phenomenal elements, the “Cartesian” view of introspection. This term is, admittedly, not quite apt, for while Descartes did hold that our introspective access to mental items *is* different in kind from our perceptual access to external objects and states of affairs, he held the further view that, insofar as a mind is a thinking thing, it has an infallible and omniscient access to its own contents. It seems, however, that any view which construes introspection in terms of direct grasping of, or acquaintance with, currently manifested phenomenal elements (and as different in kind from perception of external objects and states of affairs), is at least in the spirit of Descartes’s view, and thus merits his name.

While the Cartesian view of introspection seems plausible in certain instances, there are grounds for doubting its general applicability. This can be seen most clearly in the case of introspective access to belief states.<sup>10</sup> Beliefs are, presumably, more or less enduring psychological states, not phenomenal elements transiently manifested in consciousness. The causal activation of a belief may be accompanied by various phenomenal happenings (e.g., the “feeling” of assentingly considering the proposition believed, or a mental image of a sentence expressing that proposition),<sup>11</sup> but the belief itself is not to be identified with these happenings. Does this imply that a subject’s knowledge of her own belief states must be inferred from the phenomenal manifestations in consciousness that are associated with them? Not if we take the naturalistic or causal view of introspection suggested by Armstrong. Direct, that is, noninferential introspective knowledge of the existence of beliefs may be conditioned by prior

---

<sup>10</sup>As in Chapter 3, I restrict my attention to explicitly represented beliefs, as opposed to implicit beliefs (or, more accurately, dispositions to believe).

<sup>11</sup>See Audi (1993), p. 346.

experience with belief-associated phenomenal happenings, but it need not be inferred from such happenings. The issue is whether the subject's introspective mechanisms are operating reliably or properly in producing a (second-order) belief that the introspected belief in question does exist.<sup>12</sup>

If our view of introspection must be naturalistic, as opposed to Cartesian, in light of our introspective access to belief states, then perhaps the notion of introspective access to causal connections is not as counter-intuitive as it initially might seem. For if a naturalistic or causal perspective can help make sense of perceptual access to causal connections in the external environment, then a similar perspective can help to make sense of introspective knowledge of causal connections. If introspection is, as Armstrong says, the "getting of information or misinformation about the current state of our minds," then there is every reason to think that we can introspectively acquire information about causal connections among belief states. So long as one's introspective mechanisms have been conditioned to reliably generate beliefs about causal relations among mental states, it is reasonable to think that one has introspective knowledge concerning such relations on many occasions.

To focus on a specific example, suppose I remark (on seeing rain out the window) that I will get wet on my way home. If asked why I believe this, I may respond "Because I forgot my umbrella this morning." Here I am obviously reporting a causal connection between my belief that I forgot my umbrella this morning and my belief that I will get wet on way home. How do I know that this causal relation obtains? One may

---

<sup>12</sup>Audi suggests that "what is not itself introspectable may be readily knowable *through* what is introspectable, as one may know one is worried through the imagery and agitation one notes in oneself. Nor need the mediation be inferential: I can know I believe *p* *through* my experience of assentingly considering the proposition that I do, even if I do not *infer* that I believe *p* from this second-order proposition. What is known *by* some mark need not be *inferred from* the premise that it has that mark" (Audi 1993, pp. 346-347). I am suggesting, contrary to Audi, that what is known "through" introspectable items may also be considered introspectable.

suggest that I merely *infer* that the one belief sustains the other because I recall having previously reasoned that I will get wet because I forgot my umbrella (or, more minimally, I recall introspecting a conscious manifestation of the latter belief followed by a conscious manifestation of the former belief). But this would not establish that my belief that I forgot my umbrella *now* sustains my belief that I will get wet on my way home. If I must infer that the one belief now sustains the other, I require an additional premise that, generally speaking, when I infer one belief from another, for at least some time thereafter I continue to hold the inferred belief on the basis of its evidentiary belief. Now, however, the manner in which I know that my belief that I will get wet on my way home is sustained by my belief that I forgot my umbrella this morning is beginning to look implausibly complicated. It is more plausible to think that my introspective faculty is conditioned to noninferentially recognize cases of causal sustenance among beliefs, and that, in the case at hand, my introspective faculty directly produces the knowledge that my belief that I will get wet on the way home is sustained by my belief that I forgot my umbrella this morning. One could protest that I know that the latter sustains the former only if I know that certain counter-factual conditionals are true, e.g., that if I did not believe that I forgot my umbrella this morning, then I would not believe that I will get wet on my way home—how, it might be asked, could I know *that* introspectively? The answer is that I may not know it introspectively, but I infer it from something I do know introspectively, namely, that my belief that I will get wet on my way home is sustained by my belief that I forgot my umbrella this morning.

To sum up, if we adopt a naturalistic, as opposed to what I have termed a “Cartesian,” view of introspective knowledge, then it becomes plausible to think we have, at times, direct introspective access to causal connections. I turn now to consider the second objection, where we will find that a similar move, i.e., a shift from a Cartesian

to a naturalistic perspective, can help to make sense of our reflective access to objectively truth-linked relations of evidential support.

*Objection 2.* There is no plausible conception of objectively truth-linked relations of evidential support on which such relations would be reflectively accessible.

Whatever the character of evidential support relations, they are, as we have seen, “objectively truth-linked,” i.e., their correct analysis must reflect the fact that there is *some* important intimate relationship between a proposition’s being evidentially supported and its being true. *Prima facie*, the most plausible way to understand this link to truth is in terms of objective probability; in at least a broad range of cases, the degree of evidential support provided a conclusion by its premise set is closely tied to or determined by the objective probability of the conclusion’s being true, given the truth of its premises. We have considered two conceptions of objective probability, the logical and the nomological conceptions, and it was suggested that either could supply the appropriate notion with which to interpret the truth-linked character of evidential support.

But some would object that from the standpoint of an internalist theory of justification, only one of these accounts of objective probability, specifically, the logical conception, could be used to explicate evidential support. Their argument might run as follows. Consider first entailment relations. These are in many instances accessible on reflection, seemingly because they are *logical* relations which hold (or fail to hold) among propositions solely in virtue of their contents. At least for simple entailment relationships (such as the relation which supports the inference from the premises that all humans are mortal and Socrates is human to the conclusion that Socrates is mortal), it seems that a subject cannot fully understand the premises without seeing or grasping or apprehending that they imply the truth of the conclusion.<sup>13</sup> It seems reasonable to think that a similar

---

<sup>13</sup>Whether or not this is the correct account of our reflective access to entailment relations, it is the account favored by many internalists, including Laurence Bonjour. The

account of the reflective accessibility of evidential relations generally is required. To say that evidential support relations are accessible on reflection is to say that in a given instance where a premise, or set of premises, evidentially supports some conclusion (while perhaps falling short of entailing it), one cannot fully understand the premises or evidentiary propositions without seeing or grasping or apprehending that they support the conclusion in question or render it probably true (in the absence of evidence to the contrary). But this in turn seems to imply, if one sticks to the analogy with entailment, that the evidential support relation is a function solely of the contents of the propositions involved, and is thus in a broad sense a logical relation which holds (or fails to hold) necessarily. In that case, the objection runs, there is no *plausible* conception of objectively truth linked relations of evidential support on which such relations would be reflectively accessible. For, as some epistemologists would argue, the prospect of logical relations of evidential support or conditional probability is highly dubious (“prone to paradox at best, perhaps mythical at worst,” as Susan Haack put it).<sup>14</sup> Accordingly, they would suggest, any theory committed to their existence is equally doubtful.

*Reply.* While I do not endorse the thoroughgoing rejection of logical probability which some theorists advocate. I have no intention here of taking up the daunting project of defending this notion. In responding to the present objection, then, I will concentrate on the much less ambitious task of arguing that internalist theories of justification, and in particular, my version of internalism, are not committed to the worrisome logical conception. The problem is to find some way of answering the reasoning which underlies the current objection--specifically, that since the nomological conception of evidential support as intimately tied to nomologically objective probability is the only

---

talk of “seeing” or “grasping” or “apprehending” entailment relations is a specific application of the view of *a priori* knowledge he defends in *BonJour* (1985), and develops much more fully in *BonJour* (1997).

<sup>14</sup>Haack (1993), p. 84. See van Fraassen (1989), ch. 12, for further criticisms of the idea of logical probability.

plausible conception, and since nomological facts and relations are not reflectively accessible (there can be no direct grasp or apprehension of such facts and relations), it is not plausible to think that evidential support relations are reflectively accessible. That evidential support is not reflectively accessible, because nomological, has been stated most explicitly by Alston: “[A] ground, G, of a belief B [is] adequate if and only if it is sufficiently indicative of the truth of B. And this being the case, it seems clear that for me to be justified in believing G to be an adequate basis for belief, B, I must have sufficient *reasons* for supposing that this truth-indication relation does hold.”<sup>15</sup> By “reasons” for supposing that the truth-indicateness relation obtains, Alston means further beliefs from which one infers that G evidentially supports or is sufficiently indicative of the truth of B. Thus, one cannot, on his view, tell on *reflection* that G is supportive or indicative of the truth of B. Since nomological relations, it seems, can only be known inferentially (one could not, for example, apprehend the relation embodied in Newton’s Law of Universal Gravitation by pure reflection), this fact must apply to evidential relations if they too are nomological.

This argument, however, presupposes a particular view of the nature of reflective knowledge. Specifically, it seems to presuppose the analog, in the case of access to evidential support relations, of the Cartesian view of introspection discussed in connection with the first objection. Recall that in response to the charge that causal relations are not introspectively accessible, I argued that they only appear to be thus inaccessible from a point of view which regards introspection as the direct grasping of phenomenal elements currently manifested in consciousness, and thus involves a kind of intimate or privileged access to these items which is different in kind from our perceptual access to external objects and states of affairs. I suggested that if one adopts a different

---

<sup>15</sup>Alston (1989), pp. 240-241. Recall that, though not explicitly asserted here, Alston believes that the “truth-indicateness” of adequate grounds or evidence must be understood in nomological terms.

view of introspection. taking a causal or naturalistic account of perceptual knowledge as the model, then it becomes plausible to think that human beings have introspective access to causal connections.

Objection 2 presupposes a Cartesian view, not of introspection, but of reflective access to evidential support relations, and of access, more generally, to logical truths and relations. Consider again our reflective access to logical relationships such as the simple entailment discussed above. Here it does *seem* as if one is directly grasping or apprehending a necessary connection among propositions in a manner quite different in kind from other modes of knowing. And if such grasping of necessary connections among propositions, conceived as inseparable from or attendant upon the understanding of those propositions, were the only sort of reflective access one could have to evidential connections, then it would indeed appear that reflective knowledge of evidential relations is impossible if such relations are not themselves logical in nature. The impossibility derives not simply from the fact that nomological relations are not necessary, and thus not logical in a strict sense. The problem is that nomological connections do not seem to be the *sort* of thing which could be grasped or apprehended in the intimate way in which entailment relations seem to be apprehended.

But the Cartesian conception of reflective knowledge which underlies the second objection is not the only the view to take of such matters, even in the case of our access to the simplest logical relations. There is a possible "naturalistic" account of such knowledge which in turn could be applied in making sense of reflective knowledge of nomological evidential support relations. For, it may be that there are certain very general truths (or even "approximate truths," as we shall see) which for the survival and evolution of our species it has been advantageous for us to find self-evident. After all, given that we *do* regard certain propositions as self-evidently true, it would surely be advantageous for our survival if the things which we regard as self-evidently true are

indeed true, or at least approximately true. In that case, the best explanation of the fact that we find some connection among propositions intuitively self-evident need *not* be that this connection is entirely determined by the contents of the propositions involved. Similarly, the best explanation of the fact that we regard some proposition as self-evidently true need not be that it is genuinely necessarily true, or even more than merely approximately true. As Robert Nozick has suggested, it may be that some non-Euclidean geometry is literally true of our physical space, but that it is vastly more advantageous for us to find the axioms of Euclidean geometry to be intuitively self-evident. Thus, in order to explain why certain truths or relationships

seem self-evident to us, one need not invoke their necessity. It might be enough that they are true, even if only contingently, even just "true enough"... and that they have held true durably for long enough to leave an imprint upon our evolutionary endowment.... The strength and depth of our intuitions about certain statements cannot be used as powerful evidence for their necessity if those statements are of a kind that, were they contingent facts, would have led to selection favoring strong intuitions of their self-evidence.<sup>16</sup>

Such a perspective does not undermine the possibility of one's having knowledge of logical (or seemingly logical) relations and truths. For if one's faculty of rational insight (i.e., the faculty which yields belief in putatively self-evident truths) functions reliably or properly, because shaped by natural selection, then one may well noninferentially know that a certain set of propositions entails some further proposition, or that a certain proposition (say that nothing can be red all over and green all over at the same time) is true, though one does not know this through a special or intimate grasp of necessity which is radically different from other ways of gaining knowledge about the world.

This naturalistic account of rational insight can provide a way to make sense of reflective access to nomological relations of evidential support. As Nozick explains:

If, frequently enough, samples of a certain sort resembled their populations,

---

<sup>16</sup>Nozick (1993), pp. 110-111.

then generalizing from samples to populations, or to the next encountered member, would frequently yield truths; and beings to whom such inferences seemed obvious and self-evident would frequently arrive at those truths.... [T]hose to whom [this] connection seems closer to evident learn it faster, thereby gaining a selective advantage, and they leave offspring distributed around their own degree of finding it evident. Over generations, then, there can be movement toward finding that connection more and more self-evident.<sup>17</sup>

The example Nozick discusses involves enumerative induction, but the account need not be limited to cases of this sort. It may, for instance, be a fact about our world that explanations which are simpler according to some measure are more likely to be true than those which are more complex, and it may be that humans tend naturally to rate potential explanations as more or less plausible or reasonable in a way which tracks their degree of simplicity or complexity on the measure in question. If we hear a crash and, upon entering the kitchen, see the cat lapping up the spilled milk, we tend naturally to regard certain explanations of this observation as intuitively more reasonable than others. We would regard it as almost self-evident that the explanation that the cat knocked over the milk is much more strongly supported than, say, that the next door neighbor covertly entered the house, spilled the milk, and slipped out the back door. It may be that the former explanation has features which render it more likely to be true in *our* world, and our ability to judge the explanation's plausibility or degree of support has been honed by natural selection to approximate that objective likelihood. We may very well *know* that the explanation in question is better supported by the evidence than certain other explanations, because that knowledge is produced by a properly or reliably functioning reflective mechanism (our faculty of rational insight), despite the fact that the evidential support relation in this case is nomological, not logical.

To summarize: In response to the first objection, that causal relations are not reflectively accessible, I suggested that they only appear to be inaccessible from within the Cartesian conception of introspection. Adopt a naturalistic view, which sees

---

<sup>17</sup>Ibid., p. 109.

introspection as grounded in a reliable, noninferential, belief producing mechanism, and it seems quite plausible that we have reflective access to causal connections. Similarly, it is only on a Cartesian conception of reflective access to logical and evidential relations that objectively truth-linked relations of evidential support appear not to be reflectively accessible. From a naturalistic perspective, which regards insight into logical and evidential relations as deriving from a reliable, noninferential, belief producing faculty, there is no difficulty in thinking of certain nomological relationships as “self-evident” in a way which allows us to have reflective knowledge of them.

This is not to say that the naturalistic account of rational insight into evidential relations which we have discussed here is *correct* necessarily. Future research could vindicate the logical conception of probability, in which case the Cartesian account of rational insight would turn out to be the most appropriate. The point is only that the internalist is not *committed* to a logical conception of evidential support relations, but could adopt a nomological conception of such relations, provided that a naturalistic account of our access to such relations is adopted.

The question we must now consider is this: While it may very well be true that causal and evidential relations are reflectively accessible on a naturalistic account, this fact cannot help in responding to the present objections if such a view is incompatible with internalist conceptions of epistemic justification. Some have indeed argued that internalism is committed to a Cartesian view of our reflective faculties, an issue to which we will now turn.

*Objection 3.* Internalism is committed to an unacceptable Cartesian account of our reflective faculties.

Hillary Kornblith has argued that internalism is committed to a “Cartesian claim of incorrigible access,” but his argument suggests that it is not incorrigibility to which he thinks internalism is committed, but the broader view that reflective access to mental

states and their interrelations is different in kind from, and not capable of being explained in the same terms as, knowledge of the external world:

The fact is that once one gives up the claim of incorrigible access to one's mental states, even an agent's mental states are at a certain epistemic remove from her. Like the "external" world, an agent's mental states are external to her beliefs about them. Thus, in giving up the claim of incorrigible access, [the internalist] gives up the possibility of going internal. There no longer is a mental realm to which one has the required cognitive access. Alternatively put, there is no longer a reason to deny that one has the relevant kind of access to the "external" world, for the epistemological contrast between internal and external evaporates. One simply has reasonably good but not perfect access to all the facts there are. Some of them are about oneself and others are not, but there is no difference here epistemologically. On my view, then, internalism would be properly motivated only if one accepted the Cartesian claim of incorrigible access.<sup>18</sup>

Kornblith maintains that internalism depends for its viability or motivation on an "epistemological contrast" between the internal realm of the subject's mental states, and the external world of objects and events around her, and that once mental states (and their features and interrelations) are seen to be "at a certain epistemic remove" from the subject the required epistemological contrast "evaporates." But it seems clear that the existence of an epistemological contrast between knowledge of internal and knowledge of external states of affairs does not depend on the former being infallible or incorrigible. Surely if our reflective knowledge of mental states (and their causal and evidential interrelations) required explication in terms entirely different from the explication of our knowledge of external states of affairs, if these two sorts of knowledge required very different sorts of theory, then there would be an epistemological contrast between the internal and the external. From the perspective of the "Cartesian" views discussed in connection with the first two objections, our reflective knowledge of causal and evidential relationships looks quite different from how it would look on the naturalistic views that were also discussed.

---

<sup>18</sup>Kornblith (1988), p. 325. The paper in which the passage occurs is primarily directed against the internalism of Laurence Bonjour, and it is Bonjour's name which in fact occurs in place of the bracketed phrase, 'the internalist'. It is clear, however, that Kornblith intends his point to be applicable to internalist views generally.

This raises the possibility that Kornblith has not quite properly framed the issue he is pressing. The question is whether internalism is committed to a view of reflective knowledge as involving a special or intimate grasp of mental states and their interrelations which is not capable of being explicated in naturalistic terms, not whether such knowledge must be held to be infallible or incorrigible.

*Reply.* Let us begin by asking what sort of view *would* be committed to the indicated epistemological contrast between the internal and the external. Intuitively, if reflective faculties are judged to be fallible, then reflective knowledge suffers a certain loss of security. Beliefs produced by reflection are not quite as trustworthy as they would be were we confident that no mistake is possible. The epistemological contrast which Kornblith thinks evaporates once the fallibility of reflective processes is acknowledged is a contrast of epistemological security. Beliefs about the internal, Kornblith seems to suggest, are no more nor less trustworthy or susceptible to error than beliefs about the external world. Thus, a view which would be most obviously committed to the epistemological contrast of which Kornblith speaks would be one which holds that beliefs about the external world are in themselves less secure or trustworthy than beliefs about internal states, and in consequence warrant trust or acceptance only if they are inferentially justified on the basis of the subject's knowledge of her own internal states. The latter must provide the "foundation" from which knowledge of the external world is derived. Thus, it is foundationalist views of this sort--i.e., of the sort which hold that beliefs about the external must be inferentially justified by beliefs about the internal--which are most clearly committed to the existence of a contrast of inherent security or trustworthiness between beliefs about internal states on the one hand, and external states on the other. But once it is seen that my version of internalism is in two important respects relevantly different from such foundationalist views, it will be equally clear that

my view is not committed to the epistemological contrast in question: so that there is no bar, on my account, to explicating reflective knowledge in naturalistic terms.

First, on my view, knowledge of one's own internal states does not play a foundational role--that is, it does not play the role of conferring justification on beliefs about the external world. Consider once again my belief that *The Sleepwalkers* is on the desk in the next room. This belief about the external world is justified by *further* beliefs about the external world (to the effect that nothing has caused the book to change its location). Near as I can tell, there is no way to infer this belief from beliefs about my own internal states. Moreover, the knowledge, or rather, *potential* knowledge, which the Access Condition requires me to have concerning this belief and the relations it bears to its evidence is not *part* of that evidence, nor does it provide premises from which that belief might be inferred. In short, I do not infer my belief that *The Sleepwalkers* is on the desk in the next room from my beliefs that it is sustained and supported by its evidence; I infer this belief from the evidence itself.

Second, what motivates the foundationalist views discussed just now is the idea that having a justified belief about the external world requires that one be in a position to respond to skepticism concerning knowledge of the external world. On these views, a belief about the external world is justified only if one possesses adequate nonquestion-begging reasons for it--that is, reasons which do not depend on beliefs about the external world. But on my view, having an inferentially justified belief about the external world is not a matter of being able to respond to this sort of skepticism but, as I tried to explain in the first chapter, of being in a secure position with respect to the truth. The desire to be in such a secure position ultimately arises from the fact that we are social beings who depend on each other for information, which brings with it the expectation that we be able in appropriate circumstances to respond to questions about our beliefs--not, however, the sort of questions that a skeptic would ask, but rather the ordinary requests for reasons or

evidence which arise during the course of everyday human affairs. There is a two-fold benefit of being in such a position. First, one fulfills one's obligations to one's fellows, and second, one runs a reduced risk of being forced under questioning to give up a belief which is in fact *true*, but for which one has failed to keep track of evidence.

Think again of my belief that *The Sleepwalkers* is on the desk in the next room. What sorts of considerations or evidence would need to be available to me in order that I be in a secure position with respect to the truth of this belief? My primary reason for thinking the belief true is that I remember having left the book on the desk a short time ago, together with collateral reasons to the effect that nothing has caused the book to change location. But what further sorts of reasons must I cite in this connection? I might cite the fact that I am the only person who has been in the house all afternoon, that no earthquake has occurred which may have shaken the book from the desk, that the book is too heavy for the cat to move, and so on. That these various contingencies have not occurred is, as it were, presupposed in asserting the fact that *The Sleepwalkers* is on the desk in the next room. There are, of course, an indefinite number of other presuppositions, each representing a way in which the book might fail to stay in the location in which it was left. One way to challenge my belief would be to raise questions about the further presuppositions. What reason do I have for thinking, it might be asked, that an extraterrestrial being has not just now materialized into the other room out of thin air and absconded with *The Sleepwalkers*? What reason do I have for thinking that material objects do not from time to time randomly wink out of existence, or move suddenly three feet to the north? One might request my reasons for thinking that there is a book called *The Sleepwalkers*, that I have ever owned this book, or that there are such things as books at all. One could ask my reason for thinking that there is a desk in the next room, that I am currently in the house I take myself to be in, or that there are such things as houses. One might challenge the presupposition that I am not a disembodied

brain in a vat whose entire life and world have been fabricated by manipulative and deceiving brain scientists. In order to be in a position to answer a skeptic concerning knowledge of the external world, I must have positive evidence for my belief that *The Sleepwalkers* is on the desk in the next room which is sufficiently strong as to definitively rule out each of these possible ways in which my belief might turn out to be false. Intuitively, however, the type and amount of evidence which ordinary human affairs require us to have for our beliefs is not so great as to rule out each of these contingencies. In order to be in a secure position with respect to the truth of my belief that *The Sleepwalkers* is on the desk in the next room, I need not be able to rule out each possible way in which it might be false. I only require adequate evidence that the sorts of ways which are most typically encountered during the course of everyday events in my surroundings do not obtain. So, again, being in a secure position with respect to the truth entails satisfying less stringent standards than are involved in responding to skepticism concerning knowledge of the external world.

Because being in a secure position requires satisfying less stringent standards of epistemic evaluation than are involved in answering a skeptic, the possibility of being in a secure position does not require that one's knowledge of one's internal states be more trustworthy than one's knowledge of the external world, or that the latter be inferentially justified on the basis of the former. As I have said, my belief that *The Sleepwalkers* is on the desk in the next room is justified by reference to further beliefs about the external world, which are themselves justified by still further beliefs about the external world. There is nothing inherently insecure in such a position.<sup>19</sup>

---

<sup>19</sup>It is unclear how the various strands of this regress of justification eventually terminate. Some may ultimately stem from noninferentially justified observational or memorial beliefs, others may form complicated closed curves. This is a matter for future research.

There is one further point I want to make before moving on. My claim that being in a secure position with respect to the truth is not the same as being in a position to answer skepticism is not *itself* a skeptical conclusion. I do not claim that having an inferentially justified belief is a matter of being in a secure position with respect to the truth because I despair of the possibility of responding to skepticism. I merely think that a view of justification which ties it to the needs and concerns of ordinary human beings as they conduct their intellectual and practical endeavors, makes sense. I think it a worthwhile goal to try to spell out in precise terms the pretheoretic desideratum which guides our cognitions, namely, being in a secure position with respect to the truth. Providing a cogent, nonquestion-begging response to skepticism concerning knowledge of the external world may still be an important and viable epistemological project. But an ordinary person's having knowledge does not require that she herself take up this project.

The main question with which I have been concerned in this section is whether the Access Condition component of (IJ) is satisfiable in principle. I do not claim to have definitively answered this question in the affirmative. For the most part, I have occupied myself with the less ambitious task of arguing against those who would suggest that the Access Condition is not satisfiable in principle because it must be understood in terms of a Cartesian account of reflective knowledge. To this end, I have tried to establish two modest theses: first, that if the proper account of reflective knowledge is to be given in naturalistic terms, then the Access Condition is satisfiable in principle, and second, that there is no incompatibility between the secure position rationale of my Access Condition and the suggested naturalistic account. I have also tried to give some reason for thinking that the best account of reflective knowledge might turn out to be the naturalistic one, which warrants to some degree an at least provisional confidence that further research will prove the Access Condition to be satisfiable in principle. Unfortunately, however, a

truly definitive answer to this question requires much further investigation into the nature of reflective knowledge.

I must now turn to an equally pressing question, whether, assuming that my account is satisfiable in principle, it is widely satisfied in practice.

### 3. *Is The Account Satisfied In Practice?*

In this section, I will consider the objection that the Basis and Access Conditions, while satisfiable in principle, impose far too stringent and unrealistic standards of justification upon ordinary human beliefs. The objection maintains that these conditions are too difficult to satisfy, given certain empirical facts about human psychology, namely, that our long term memories lack sufficient storage capacity and our reflective faculties are too weak to allow satisfaction of my proposed conditions in normal circumstances. Since, it is claimed, most of our beliefs are justified, and since a great many of these must be inferentially justified if justified at all, the Basis and Access Conditions cannot be genuinely necessary for inferential justification. I will discuss this objection first in application to the Basis Condition, then in application to the Access Condition. I will argue that while my conditions may not be widely satisfied, they serve a useful function as slightly idealized standards of inferential justification which, as I have tried to argue, give precise form to a pretheoretic cognitive goal that we strive to obtain.

The issue I will be focusing on in the following discussion is whether ordinary *adult* human beings are able to satisfy my conditions of inferential justification. It is clear that most young children, insofar as they lack workable concepts of causal sustenance and evidential support, are unable to satisfy the Access Condition and therefore, on my view, lack inferentially justified beliefs. This is not to say that they fail to have justified beliefs altogether. They may well have *noninferentially* justified beliefs (e.g., perceptual, introspective, and memorial beliefs), and they may have inferentially justified beliefs in a

derivative sense, i.e., beliefs for which they possess further evidentiary beliefs, and which in the fullness of time will be inferentially justified by those further beliefs. I think that such a stance with regard to the beliefs of children is warranted by the fact that the notion of the secure position seems most comfortably applied to the beliefs of adults, and in as much as we wish to evaluate individuals' with respect to the secureness of their cognitive positions, this sort of evaluation will be most naturally directed at adults (or more specifically, at individuals who possess the requisite concepts of evidential support and causal sustenance). Thus, in my view, the salient question is whether ordinary adults can in fact satisfy my conditions of inferential justification. I shall not accept it as a damaging consequence that my account prevents young children (and, for that matter, the more intelligent animals) from having inferentially justified beliefs in the strict sense.

*Objection 4.* The Basis Condition is not psychologically realistic. Very few of a typical adult's beliefs satisfy this condition.

The Basis Condition requires inferentially justified beliefs to be sustained by evidentiary beliefs retained in long term memory. The problem is that for a great many beliefs which would appear to be inferentially justified if justified at all, no specific evidence is retained in long term memory which could provide the required causal sustenance. Some examples of this sort we have discussed already (in Chapter 2). I may believe that a certain friend is dishonest, but I may have failed to retain specific instances of his dishonest behavior. I believe that Thomas Jefferson founded the Democratic Party, that the Special Theory of Relativity was published in 1905, that Istanbul was once named 'Constantinople.' It seems that I have no genuine evidence for these beliefs. Their continued existence is sustained by a vague sense that at some time and place, I once witnessed, or heard of, or read of, their truth. But memories of these sorts do not seem capable of providing much evidence.

A striking example of the failure of human beings to retain specific evidence for their beliefs can be found in laboratory studies of the phenomenon of belief perseverance. In one such study, subjects were given a task involving social sensitivity skills (they were asked to sort authentic from fabricated suicide notes). The subjects were then told whether they performed at an above average, average, or below average level, but their performance rating was arbitrarily chosen by the experimenters and had no relationship to their actual performance of the task. Shortly afterwards, subjects were given a thorough “debriefing” in which they were told that their performance level had been arbitrarily chosen by the experimenters, and that they accordingly had been given no real evidence as to their skill in executing the task in question. Nevertheless, when asked to predict their performance in similar tasks in the future, those who had been given a high rating predicted that they would perform at an above average level, while those who had been given lower ratings predicted that they would perform at a less than average level. The beliefs concerning their social sensitivity skills which the subjects derived from their performance in the studies persevered even after the putative evidence had been thoroughly discredited.<sup>20</sup>

There are a number of possible explanations for the belief perseverance phenomenon, for example, that the beliefs in question are like bad habits of which, try as they might, the subjects are unable to rid themselves, or that the subjects simply fail to notice that the evidence for their beliefs has been undermined. But Gilbert Harman has argued that these explanations are unsatisfactory. First, the subjects do not act as if their beliefs are bad habits, and seem perfectly comfortable in holding them. And second, he claims, it “stretches credulity to suppose that people always keep track of the sources of their beliefs but often fail to notice when these sources have been undermined.”<sup>21</sup> Thus,

---

<sup>20</sup>This and other studies are described in Nisbett and Ross (1980), chap. 8.

<sup>21</sup>Harman (1986), p. 41.

according to Harman, the best explanation of the belief perseverance phenomenon is that people in general do not retain evidence for their beliefs. Once a belief has been formed, they see no reason to preserve evidence for it, and thus are unable to see when, subsequently, their evidence has been discredited. On Harman's view, the beliefs involved in the perseverance studies are not unlike my beliefs discussed a moment ago (that my friend is dishonest, that Thomas Jefferson founded the Democratic Party, and so on), in that none seem supported by specific evidence, but continue to exist through a kind of "cognitive inertia." Once formed, they continue to exist in the absence of the initial evidential impetus which created them.<sup>22</sup>

A closely related difficulty for the Basis Condition stems from what may be termed the "cognitive division of labor." It is an obvious fact of life that a great many of the things we believe depend upon information we receive from other individuals, especially experts or authorities, whether in person, or in books, or in the media. To the extent that our beliefs depend on expert testimony, we will not be ourselves in possession of the evidence which supports our beliefs, and indeed may be unable to grasp the evidence should the belief in question pertain to some abstruse field of knowledge such as quantum mechanics or relativistic gravitational physics. Even on such a simple matter as my belief that people have walked on the Moon, I do not understand in any great detail the process involved in sending people to the Moon, and I did not go along for the ride.<sup>23</sup> It seems that the most I can do by way of giving evidence for this belief is to cite the reliability of the authorities from whom I learned that the endeavor took place. But this

---

<sup>22</sup>Nisbett and Ross, as well as Goldman (1986), offer an explanation of belief perseverance which Harman does not discuss, namely, that the subjects find further beliefs in long term memory which cohere with or corroborate the beliefs concerning their social sensitivity skills. These further beliefs serve to sustain the beliefs concerning their sensitivity once the original evidence is undermined. But for the purposes of the present discussion, I will not pursue this explanation.

<sup>23</sup>This example is due to Bill Talbott who, in personal communication, has pressed the difficulties I am here considering.

itself is rather thin evidence, particularly when one considers that I may not have any specific evidence for my belief that the authorities in question are reliable (and I may have only a vague memory as to who these authorities are).<sup>24</sup>

The problem that we have been discussing in connection with Objection 4 can be summed up as follows. The Basis Condition requires inferentially justified beliefs to be sustained by adequate evidence retained in long term memory. But in fact a great many apparently inferentially justified beliefs, such as my belief that people have walked on the Moon, or that my friend is dishonest, or that Thomas Jefferson founded the Democratic Party, are not sustained by anything which deserves to be called *adequate* evidence. The most that I possess by way of evidence for these beliefs might be called "reliability reasons." I may claim to remember that my friend has perpetrated many deceptions, or that I once read that Thomas Jefferson founded the Democratic Party, and I can claim that my memory is in general reliable in such matters. But it is unclear to what extent I can give evidence as to the reliability of my memory. Similarly, I can claim that my beliefs that people have walked on the moon results from reliable authorities, though I may be incapable of providing evidence for their reliability. It is clear, however, that such reliability reasons are not sufficient to satisfy the Basis Condition if I lack genuine evidence of my own, or others, reliability in such matters.

*Reply.* It seems clear that if the beliefs we have discussed are indeed justified, then they must be inferentially justified. I seem capable of having memory *impressions* only of singular events in my own life history. The scope of beliefs which could be noninferentially justified on their basis therefore seems limited to beliefs concerning the singular events themselves, which excludes the beliefs we considered above. It would, for example, be impossible for me to have a memory impression of Thomas Jefferson's

---

<sup>24</sup>Alston has argued for this point. "How many of us," he asks, "can, on reflection, come up with adequate evidence on which to base the belief that a given putative authority is to be relied on? Very few of us" (Alston 1989, p. 241).

founding the Democratic Party. The most I could be capable of having is a memory impression of having read or heard of Thomas Jefferson's founding the Democratic Party, but I have in fact retained no such specific impression. Furthermore, my belief that people have walked on the Moon, and a great many other beliefs I hold as a result of others' testimony, are not the sort of beliefs which could be noninferentially justified by *any* means.<sup>25</sup>

In addition, it seems clear that all of these beliefs fail strictly speaking to satisfy the Basis Condition. The problem is not that some or all of these beliefs are based on reliability reasons. For there is no reason in principle why a belief could not be based (in part) on a reliability reason and still satisfy the Basis Condition. For example, my belief that people have walked on the Moon *could* be held on the inferential basis of beliefs that a particular source S (an expert or a textbook, say) told me that people have walked on the Moon, and that S is a reliable source on such matters (and if I have evidence for these latter beliefs they might well confer justification on the former belief). But while it is possible for a belief to satisfy the Basis Condition by being based in part on a reliability reason, the beliefs discussed above still fail to satisfy this condition. For the fact remains that I have failed to retain *adequate* evidence for these beliefs. It is not enough that I have a vague sense that my memory is reliable on questions of this sort, or that I did once learn thus and so from a reliable source (though I cannot remember exactly who it was). I see no way of escaping the conclusion that, if these beliefs are any indication, a great many of my beliefs fail to satisfy the Basis Condition. The crucial question is whether we should reject the Basis Condition in light of this.

---

<sup>25</sup>Plantinga (1993b, p. 79) has argued that at least some beliefs based on testimony must be noninferentially justified (or warranted, to use his term), since otherwise a vicious regress of testimonial beliefs would result. But his argument raises complicated issues beyond the scope of our discussion. Accordingly, I will assume that beliefs held on the basis of testimony must be inferentially justified if justified at all.

The answer, I believe, is that we ought to accept the Basis Condition, not as an incontrovertible standard of justification which beliefs *must* satisfy in order to have any epistemic value, but as a tool for making a certain important kind of evaluation of our beliefs. For while there may be a sense in which the beliefs we have discussed are justified, we must not overlook the fact that they are legitimately open to criticism. Consider, for example, my belief that my friend is dishonest. The problem is that I am *disappointed* in myself for not doing a better job in keeping track of instances of dishonest behavior on the part of my friend. After all, it is no small matter to accuse a friend of being dishonest. Is it not just a bit unreasonable for me to be so confident when I cannot remember a specific incident in which my friend was in fact dishonest? Notice that my belief that my friend is dishonest may well be reliably sustained by certain memorial processes, and thus is (in an objective sense) very probably true, and so can be (in *some* sense) “trusted.” But since I myself have done a bad job of retaining evidence for my belief, it is still open to criticism in an obvious respect. There is an important dimension of epistemic evaluation along which it fares rather poorly, despite the fact that, from another perspective, my belief has the admirable property of being reliably sustained and thus very probably true. And it seems to me that I myself, *qua* “ordinary person,” can recognize this sense in which my belief is deficient. Indeed, *all* of the beliefs discussed above are defective precisely because I am not in a secure position with respect to their truth.

Now, it could reasonably be pointed out that it is not possible for me to be in a secure position with respect to the truth of *all* of my beliefs, since I am *unable* to retain the required evidence. The problem is not that I am lazy or forgetful (although of course I am). The problem is that I just do not have enough storage capacity to satisfy the Basis Condition with respect to all of my beliefs (or rather, all of my beliefs that are not justified in some other way). I could not do it, no matter how hard I tried. This does not

show, however, that the goal of being in a secure position with respect to the truth is a pointless one. It suggests rather that I ought to, as it were, conserve my efforts so that I will at least be in a secure position with respect to the truth of those beliefs which are most important to me. I will, in light of my limitations, put more emphasis on being in a secure position in some situations than in others. For instance, the fact that I am not in a secure position with respect to the truth of my belief that my friend is dishonest causes me more concern than the fact that I am not in such a position with respect to the truth of my belief that Thomas Jefferson founded the Democratic party. Since nothing much hangs on the latter belief, its being reliably sustained in my memory is perhaps as much positive epistemic worth as it needs (though not as much as it could have, and in this respect it is defective). But regardless of my cognitive limitations, I feel that I ought to have retained evidence for my belief that my friend is dishonest, and in view of this fact, it seems inappropriate to regard it as justified simply because it satisfies some less stringent standard of epistemic evaluation.

Thus, the fact that a fair number of our beliefs fail to satisfy the Basis Condition should not be thought to show that this condition is a useless tool of epistemic evaluation. Rather, it means that the Basis Condition should be thought of as a slightly idealized standard of justification (in the sense that it does not purport to reflect our actual cognitive efforts in all situations) which is perhaps best applied in those circumstances in which we strive to be in a secure position with respect to the truth. Something similar may be said about the Access Condition, as we will see in considering the next objection.

*Objection 5.* The Access Condition is not psychologically realistic. Very few of a typical adult's belief satisfy this condition.

If the Basis Condition is not realistic from the standpoint of the psychology of ordinary human beings, the Access Condition might be thought even more dubious from this perspective. In the first place, in those cases where the Basis Condition is not

satisfied, the Access Condition will typically not be satisfied as well. After all, if no specific evidence for a belief is retained, no such evidence will be accessed on reflection. But even in those cases where a belief is correctly inferentially based on adequate evidence, some will claim that it is doubtful that the subject could be expected to access that evidence on reflection and tell *that* it sustains and supports her belief. The doubtfulness of this prospect stems from two sources. First, there is empirical data which demonstrates the weakness of human reflective faculties in at least certain settings. To give one example, subjects in a certain study were asked to compare four exactly similar items of clothing laid out on a table at a shopping mall. Their task was to decide which of these garments was of the highest quality, and in doing so they tended strongly to prefer the garment which was situated on their right. Afterward, they insisted that the position of the items played no role in their assessment, indicating an inability on the part of these subjects to reflectively uncover the causal factors underlying their belief. These results, as well as those of other studies in a similar vein, suggest that in at least some circumstances human beings lack introspective access to the actual cognitive factors which influence belief formation.<sup>26</sup> Indeed, some have drawn a much more pessimistic conclusion from the studies in question:

People often cannot report accurately on the effects of particular stimuli on higher order, inference-based responses.... The accuracy of subjective reports is so poor as to suggest that any introspective access that may exist is not sufficient to produce generally correct or reliable reports.... When reporting on the effects of stimuli, people may not interrogate a memory of the cognitive processes that operated on the stimuli; instead, they may base their reports on implicit, a priori theories about the causal connection between stimulus and response.... Subjective reports about higher mental processes are sometimes correct, but even the instances of correct report are not due to direct introspective awareness. Instead, they are due to the incidentally correct employment of a priori causal theories.<sup>27</sup>

---

<sup>26</sup>The study just described is from Nisbett and Wilson (1977). See also Bem and McConnell (1970) for further studies. Baergen (1995) surveys all of this data and concludes that internalism is not viable in light of it.

<sup>27</sup>Nisbett and Wilson (1977), p. 233. Quoted in Baergen (1995), p. 38.

But apart from the weak and unreliable character of human reflective faculties in some circumstances, critics of internalism will also maintain that an examination of particular cases of inferentially justified belief suggests that the Access Condition, at least as formulated here, is not in general satisfied. We have already seen, for example, that Alston (among others) claims that while a subject's evidence may itself be accessible to her, she cannot be expected to tell that the evidence sustains and supports her belief. And Bonjour maintains only that the evidence must be available to the subject *in principle*. Some might argue that if even internalists are willing to concede that an access condition as strong as the one defended here is not in general satisfied, it must be a very dubious condition indeed.

*Reply.* There are two points raised in the objection which must be addressed. The first is the general claim that empirical data indicates a pervasive weakness and unreliability in human reflective powers, indeed a weakness and unreliability of such startling proportions as to make it doubtful that the Access Condition is widely satisfied. The second is the claim that an examination of paradigm cases of inferential justification reveals that the Access Condition is often not satisfied. Let us take these in turn.

First, to offer a general observation, the inaccuracy or unreliability of introspective reports cannot be as widespread as the considerations just discussed would suggest. In our dealings with people in everyday situations, we assume that others are endowed with a reasonable degree of self-awareness. On the basis of this assumption we predict one another's behavior, and plan our own actions by reference to such predictions. This assumption is a pervasive, integral part of living together in a society. If it is false, then it is a miracle that our predictions of behavior are not routinely incorrect, and that our cooperative efforts are not sabotaged.<sup>28</sup> Moreover, the suggestion that our access to our own mental states and processes is largely inferential, and that our putative introspective

---

<sup>28</sup>Shoemaker (1988) makes this point.

reports are correct in virtue of the "incidental correctness" of certain causal theories, is mistaken. Any inference I make concerning my own mental states or processes will depend on my having introspective access to certain other mental states or processes. For example, I may notice that I am having certain thoughts and conclude, "I must be angrier than I realized." It is fairly clear that this inference depends on my having introspective access to the thoughts in question. Or, someone might ask whether I prefer chocolate ice cream to vanilla, and I might respond, "I tend to order vanilla more often than chocolate, so I guess I actually like vanilla better." Again, it appears that my answering this question depends on my having introspective access to certain memories of ordering vanilla ice cream instead of chocolate. So it seems that inference concerning one's own mental states is parasitic on one's having introspective access to certain other mental states.

It is worth noting as an aside that empirical data of the sort discussed in connection with Objection 5, purporting to show the widespread inaccuracies in our introspective reports, do not *directly* count against the Access Condition. For the condition states a requirement on inferentially justified belief, but the studies in question do not provide cases of inferentially justified beliefs in which the Access Condition is not satisfied. For instance, the subjects who believe that the garment on the right hand side of the table is superior to the others are not inferentially justified in believing this, and so their beliefs are not counter-examples to my Access Condition.<sup>29</sup> Thus, the second point in support of the fifth objection, that the Access Condition is not satisfied in paradigm

---

<sup>29</sup>I should point out that there are studies which demonstrate the reliability of introspection in at least some circumstances. For instance, Nelson and Dunlosky (1991) describe a study in which subjects were very accurate at predicting their ability to recall information, provided that a few minutes' delay had occurred between their memorizing the information in question and their being asked to estimate their ability to recall it. Moreover, such estimations would appear to be introspective in character.

cases of inferentially justified belief, constitutes the more serious threat. I shall now consider it.

It does not strike me as implausible that the Access Condition is satisfied in many of the paradigm cases of inferentially justified belief (provided that the Basis Condition is satisfied). For example, Marshall Swain, who does *not* accept an access condition, has offered this as a paradigm case of inferential justification:

Throughout her childhood Roberta has believed that whales are fish, largely because she knows they live in water and swim like fish. While studying her high school biology book, however, she reads that whales are actually mammals and not fish. Somewhat amazed by this news, she asks her biology teacher if there isn't perhaps a mistake in the text. Her teacher assures her that the text is correct. Over the weekend she goes to the local sea aquarium and studies the whales there, noting certain mammalian characteristics for herself. On the basis of the various beliefs that she has thus acquired, she comes to believe that all whales are mammals. Her belief in this generalization on the basis of her reasons is, I take it, epistemically justified.<sup>30</sup>

It strikes me as quite reasonable to think that Roberta satisfies the Access Condition in holding her belief that all whales are mammals. If asked why she believes that all whales are mammals, she could respond (though perhaps in less sophisticated language), "My biology textbook asserts that whales are mammals, and it is unlikely that a book with misleading information on such a basic topic would be accepted by my school district. Moreover, my teacher confirms that whales are mammals, and as a teacher of biology, if she confidently makes such a claim, it is very likely true. Also, I myself have studied whales at the aquarium and have observed that they have defining mammalian characteristics (for example, the mothers nurse their young). I know that all of this information gives me a pretty good reason to think that all whales are mammals." Of course, as a high school student, Roberta would probably not put her reasons into exactly these terms, but anything close would put her in a fairly secure position with respect to the truth of her belief.

---

<sup>30</sup>Swain (1981), p. 94.

It will have to be conceded, however, that cases like this are atypical in a certain respect. For Roberta is obviously concerned enough with the question whether all whales are mammals to seek and retain evidence for her belief over a period of days. She may not always be so concerned to gather evidence in this way, and the same must be said for the rest of us. We do not *always* care whether we are in a secure position with respect to the truth.

But while we do not always care whether we are in a secure position with respect to the truth, there is still a point to focusing philosophically on that desideratum in order to articulate conditions which capture it in a precise way. Otherwise we would have no way of distinguishing philosophically between Roberta's belief that all whales are mammals, and the belief of one who accepts this on the basis of a reliable source (such as a biology text) without regard for the reliability of that source, or the possibility of finding corroborating evidence for it. That these beliefs should be distinguished is surely implicit in our common sense epistemic evaluations, for there are many situations in which we attach a good deal of importance to seeking corroborating evidence for our sources, not merely taking them for granted. Reputable newspaper reporters, for example, will as a rule rely on a source only if the information he or she has offered has been confirmed in some way.

So, the fact that our beliefs fail to satisfy the Access Condition in certain circumstances, does not show that some other less stringent standard is the appropriate one to use in making epistemic evaluations. For if we adopted in all situations a less stringent standard (such as the standard of having beliefs which are produced by a reliable mechanism) then we would ignore the fact that in many situations, most conspicuously situations of deliberation and directed inquiry, but in others as well, we strive for something more, namely, to be in a secure position with respect to the truth. Thus, the appropriateness of the Access Condition as a tool of epistemic evaluation stems not from

its capturing a standard which we always satisfy, or even always try to satisfy, but rather from the fact that it represents an epistemic ideal which we strive to attain (and sometimes succeed in attaining) in many everyday situations. Like the Basis Condition, it need not reflect our cognitive efforts in all circumstances in order to explicate an important cognitive goal.

#### *4. Summary.*

In this chapter, I tried to respond to objections which allege that my account of inferential justification (embodied in the Basis and Access Conditions) is satisfiable neither in principle nor in practice. I tried to respond to objections which maintain that the account is not satisfiable in principle by arguing that my project of explicating the intuitive notion of being in a secure position with respect to the truth is compatible with a naturalistic picture of reflective knowledge according to which we have reflective access both to relations of causal sustenance and objectively truth-linked relations of evidential support among our beliefs. I tried to respond to objections which maintain that the account is not satisfiable in practice by arguing that my two conditions capture an ideal which we strive to attain in many circumstances, indeed sometimes succeed in attaining, and that therefore the cases in which we fail to satisfy these conditions should not encourage us to adopt a less stringent standard of epistemic evaluation in all situations.

## Bibliography

- Alston, William P. (1989). *Epistemic Justification*. Ithaca: Cornell University Press.
- Alston, William P. (1991). *Perceiving God: The Epistemology of Religious Experience*. Ithaca: Cornell University Press.
- Annis, David. (1978). "A Contextualist Theory of Epistemic Justification." *American Philosophical Quarterly*, 15, pp. 213-219.
- Armstrong, D. M. (1968). *A Materialist Theory of the Mind*. London: Routledge & Kegan Paul.
- Armstrong, D. M. (1973). *Belief, Truth, and Knowledge*. London: Cambridge University Press.
- Audi, Robert. (1982). "Believing and Affirming." *Mind*, 91, pp.115-120.
- Audi, Robert. (1988). *Belief, Justification, and Knowledge*. Belmont: Wadsworth.
- Audi, Robert. (1993). *The Structure of Justification*. New York: Cambridge University Press.
- Baergen, Ralph. (1995). *Contemporary Epistemology*. Orlando: Harcourt Brace.
- Bem, D. J., and H. K. McConnell. (1970). "Testing the Self-Perception Explanation of Dissonance Phenomena: On the Salience of Premanipulation Attitudes." *Journal of Personality and Social Psychology*, 14, pp. 23-31.
- BonJour, Laurence. (1985). *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- BonJour, Laurence. (1986). "A Reconsideration of the Problem of Induction." *Philosophical Topics*, 14, pp. 93-124.
- BonJour, Laurence. (1992). "A Rationalist Manifesto." *Canadian Journal of Philosophy*, supplementary volume 18, pp. 53-88.
- BonJour, Laurence. (1994). "Against Naturalized Epistemology." *Midwest Studies in Philosophy*, 19, pp. 283-300.
- BonJour, Laurence. (1995). "A Moderate Rationalism." *Philosophical Topics*, 23, pp. 47-78.
- BonJour, Laurence. (1997). *In Defense Of Pure Reason*. Forthcoming from Cambridge University Press.

- Burks, Arthur W. (1963). "On the Significance of Carnap's System of Inductive Logic for the Philosophy of Induction." *The Philosophy of Rudolph Carnap*, P. A. Schilpp, ed. La Salle: Open Court.
- Burks, Arthur W. (1977). *Chance, Cause, Reason*. Chicago: Chicago University Press.
- Carnap, Rudolph. (1950). *The Logical Foundations of Probability*. Chicago: Chicago University Press.
- Carnap, Rudolph. (1952). *The Continuum of Inductive Methods*. Chicago: Chicago University Press.
- Chisholm, Roderick. (1980). "A Version of Foundationalism." *Midwest Studies in Philosophy*, 5, pp. 543-564.
- Chisholm, Roderick. (1989). *Theory of Knowledge*, third edition. Englewood Cliffs: Prentice-Hall.
- Conee, Earl. (1988). "The Basic Nature of Epistemic Justification." *Monist*, 71, pp. 389-404.
- Conee, Earl, and Richard Feldman. (1985). "Evidentialism." *Philosophical Studies*, 48, pp. 15-34.
- Feldman, Richard. (1988a). "Epistemic Obligations." *Philosophical Perspectives*, 2, *Epistemology*, James Tomberlin, ed. Atascadero: Ridgeview.
- Feldman, Richard. (1988b). "Having Evidence." *Philosophical Analysis*, David Austin, ed. Dordrecht: Kluwer Academic Publishers.
- Feldman, Richard. (1992). "Evidence." *A Companion to Epistemology*, Jonathan Dancy and Ernest Sosa, eds. Cambridge: Blackwell.
- Firth, Roderick. (1978). "Are Epistemic Concepts Reducible to Ethical Concepts?" *Values and Morals*, A. I. Goldman and J. Kim, eds. Dordrecht: D. Reidel.
- Fumerton, Richard. (1985). *Metaphysical and Epistemological Problems of Perception*. Lincoln: University of Nebraska Press.
- Fumerton, Richard. (1988). "The Internalism/Externalism Controversy." *Philosophical Perspectives*, 2, *Epistemology*, James Tomberlin, ed. Atascadero: Ridgeview.
- Fumerton, Richard. (1995). *Metaepistemology and Skepticism*. Lanham: Rowman & Littlefield Publishers.
- Gettier, Edmund. (1963). "Is Justified True Belief Knowledge?" *Analysis*, 23, pp. 121-123.

- Goldman, Alvin. (1967). "A Causal Theory of Knowing." *Journal of Philosophy*, 64, pp. 355-372.
- Goldman, Alvin. (1979). "What is Justified Belief?" *Justification and Knowledge*, George S. Pappas, ed. Dordrecht: D. Reidel.
- Goldman, Alvin. (1986). *Epistemology and Cognition*. Cambridge: Harvard University Press.
- Goldman, Alvin. (1988). "Strong and Weak Justification." *Philosophical Perspectives*, 2, *Epistemology*, James Tomberlin, ed. Atascadero: Ridgeview.
- Goldman, Alvin. (1992). *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge: MIT Press.
- Haack, Susan. (1993). *Evidence and Inquiry*. Oxford: Blackwell.
- Harman, Gilbert. (1986). *Change in View*. Cambridge: MIT Press.
- Hume, David. (1938). *An Abstract of a Treatise of Human Nature*. Cambridge: Cambridge University Press.
- Kapitan, Thomas. (1985). "Reliability and Indirect Justification." *Monist*, 68, pp. 277-287.
- Kim, Jagwon. (1988). "What is 'Naturalized Epistemology'?" *Philosophical Perspectives*, 2, *Epistemology*, James Tomberlin, ed. Atascadero: Ridgeview.
- Kitcher, Philip. (1992). "The Naturalists Return." *Philosophical Review*, 101, pp. 53-114.
- Kornblith, Hilary. (1980). "Beyond Foundationalism and the Coherence Theory." *The Journal of philosophy*, 72, pp. 597-612.
- Kornblith, Hilary. (1983). "Justified Belief and Epistemically Responsible Action." *The Philosophical Review*, 92, pp. 33-48.
- Kornblith, Hilary. (1988). "How Internal Can You Get?" *Synthese*, 14, pp. 313-327.
- Kornblith, Hilary. (1992). "Naturalized Epistemology." *A Companion to Epistemology*, Jonathan Dancy and Ernest Sosa, eds. Cambridge: Blackwell.
- Kornblith, Hilary. (1993). *Inductive Inference and Its Natural Ground*. Cambridge: MIT Press.
- Kornblith, Hilary. (1994). "Introduction: What is Naturalistic Epistemology?" *Naturalizing Epistemology*, Hilary Kornblith, ed. Cambridge: MIT Press.

- Kornblith, Hilary. (1995). "Naturalistic Epistemology and its Critics." *Philosophical Topics*, 23, pp. 237-255.
- Kyburg, Henry E. (1965). "Comments on Salmon's 'Inductive Evidence.'" *American Philosophical Quarterly*, 2, pp. 245-264.
- Lehrer, Keith. (1974). *Knowledge*. London: Oxford University Press.
- Lehrer, Keith. (1990). *Theory of Knowledge*. Boulder: Westview Press.
- Moore, G. E. (1962). *Commonplace Book 1919-1953*. London: George Allen & Unwin
- Moser, Paul. (1989). *Knowledge and Evidence*. New York: Cambridge University Press.
- Nelson, Thomas and John Dunlosky. (1991). "When People's Judgments of Learning (JOLs) Are Extremely Accurate At Predicting Subsequent Recall: The "Delayed-JOL Effect." *Psychological Science*, 2, pp. 267-270.
- Nisbett, Richard and Lee Ross. (1980). *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs: Prentice-Hall.
- Nisbett, Richard and Timothy Wilson. (1977). "Telling More Than We Can Know: Verbal Reports on Mental Processes." *Psychological Review*, 84, pp. 231-259.
- Nozick, Robert. (1981). *Philosophical Explanations*. Cambridge: Harvard University Press.
- Nozick, Robert. (1993). *The Nature Of Rationality*. Princeton: Princeton University Press.
- Plantinga, Alvin. (1993a). *Warrant: The Current Debate*. New York: Oxford University Press.
- Plantinga, Alvin. (1993b). *Warrant and Proper Function*. New York: Oxford University Press.
- Pollock, John. (1986). *Contemporary Theories of Knowledge*. Savage: Rowan and Littlefield.
- Pollock, John. (1990). *Nomic Probability and the Foundations of Induction*. New York: Oxford University Press.
- Quine, W. V. O. (1969). *Ontological Relativity and Other Essays*. New York: Columbia University Press.

- Russell, Bertrand. (1940). *An Inquiry Into Meaning and Truth*. Baltimore: Penguin.
- Shoemaker, Sydney. (1988). "On Knowing One's Own Mind." *Philosophical Perspectives*, 2, *Epistemology*, James Tomberlin, ed. Atascadero: Ridgeview.
- Sosa, Ernest. (1993). *Knowledge in Perspective*. New York: Cambridge University Press.
- Steup, Matthias. (1996). *Contemporary Epistemology*. Upper Saddle River: Prentice-Hall.
- Swain, Marshall. (1979). "Justification and the Basis of Belief." *Justification and Knowledge*, George S. Pappas, ed. Dordrecht: D. Reidel.
- Swain, Marshall. (1981). *Reasons and Knowledge*. London: Cornell University Press.
- Swain, Marshall. (1988). "Alston's Internalistic Externalism." *Philosophical Perspectives*, 2, *Epistemology*, James Tomberlin, ed. Atascadero: Ridgeview.
- Talbott, W. J. (1990a). Preface to *The Reliability of the Cognitive Mechanism*. New York: Garland Publishing.
- Talbott, W. J. (1990b). *The Reliability of the Cognitive Mechanism*. New York: Garland Publishing.
- van Fraassen, Bas. (1989). *Laws and Symmetry*. New York: Oxford University Press.
- Winters, Barbara. (1983). "Inferring." *Philosophical Studies*, 44, pp. 201-220.

## Vita

Stephen Fogdall was born in Seattle in 1968. He has received three degrees from the University of Washington: his B.A. in 1990, his M.A. in 1992, and his Ph.D. in 1997. He enjoys his family and his friends. He likes to read books, and do things outdoors.